

LINUX JOURNAL

The Monthly Magazine of the Linux Community

Cutting Edge Linux

222454
\$449


COAS: Caldera Open Administration System

Virtual Network Computing

Hunting Hurricanes

Csound for Linux

KDE and GNOME



USA \$5.00 CAN \$6.50


www.linuxjournal.com

02



Wearable Computers

Where do you want to go tomorrow?

our hardware vendor thinks Linux is a rare tropical disease. Your software vendor thinks you like rebooting every time you encounter a “General Protection Fault.” Been there, done that. Maybe that’s where you are today, but it’s not where you want to go tomorrow.

You want a different kind of company — one that speaks your language, understands your OS, and sells your kind of systems. The first and best Linux hardware company — VA Research. Since 1993 we’ve been designing systems from the ground up to support Linux. No need to spend your time testing, benchmarking, and debugging. We’ve done that for you. Check out the

benchmarks on our Web page. You have work to do, and VA Research can get you there better and faster.

From laptops to workstations to multi-processor servers our Linux systems give you the performance and reliability you demand. Whatever your application, we can find a solution that fills your needs.


**Take your first step toward tomorrow by
calling VA Research today at 1.888.LINUX.4U
or visiting www.varesearch.com.**

Let VA Research take you there faster.

VARBook 100

Linux on the NEC LX laptop.

- 266 MHz Pentium II CPU
- 96MB RAM
- 14" 1024x768 active screen
- Up to 6GB disk
- Modem, floppy, and CD-ROM
- Ethernet
- Carrying case



"A truly robust Unix clone...all the features of commercial versions... I pulled out my credit card and ordered a (VA Research) system for my personal use."

Byte Magazine

"VA Research's VARStation ...is the enterprise-computing equivalent of speaking softly and carrying a big stick..."

InfoWorld

"The VarStation has been built from the ground up to run Linux and run it well. VA Research didn't scrimp on hardware and devoted a great deal of effort to ensure that the system is tuned to work as configured."

InfoWorld


"This is an excellent combination of hardware for Linux. It is probably the fastest single-processor x86-based system available (VA Research has some multi-processor server systems, too), and all of the equipment works with Linux and is pre-configured for it."

Linux Journal

VARStation YMP

Superior integer, floating point, and multimedia performance.

- One or Two Pentium II CPUs
- Speeds up to 450 MHz
- Up to 512MB SDRAM
- 4.5GB Ultra Wide SCSI Disk
- Ultra 2 SCSI Available
- 8MB Matrox Millenium II Video
- 100Mb/s Ethernet
- 32X CD-ROM Drive




| Model | SPECint_rate_base95 |
|-------------------------------------|---------------------|
| VA Research VARServer 4000 | 247 |
| VA Research VARStation YMP | 215 |
| Sun Ultra 2 Creator (3D) Model 2200 | 133 |
| Digital AlphaStation 500/500 | 113 |

| Model | SPECint95_base |
|---------------------------|----------------|
| VA Research VARStation 28 | 13.1 |
| Sun Ultra 60 | 13.0 |
| IBM RS/6000 43P | 5.96 |
| SGI 02 R5000SC | 4.76 |

VARServer 3100

Rackmount system.

- Two 450 MHz Pentium II CPUs
- Up to 512MB SDRAM
- Up to 45GB RAID-5 Storage
- 12GB 4mm SCSI Tape
- 32X CD-ROM Drive
- 100Mb/s Ethernet
- Watchdog System Monitor
- UPS




And...

Alpha workstations, Beowulf clusters, dual-boot systems, rackmount cases, OEM products, and full custom systems.

VARServer 4000

Mission-critical RAID serving.

- Four Pentium Pro 200 MHz CPUs
- Up to 4GB ECC RAM
- Up to 200GB RAID-5 storage
- DPT Raid Controller
- 20GB native DLT tape
- 100Mbps Ethernet
- Watchdog System Monitor
- 2GB Boot Disk
- Redundant Power Supplies
- UPS



VA Research, Inc.
1235 Pear Ave., Suite 109
Mountain View, CA 94043
1.888.LINUX.4U
Tel: +1.650.934.3666
Fax: +1.650.964.7668

email: sales@varesearch.com
web: http://www.varesearch.com

1.888.LINUX.4U

President* Declares FeBRUary International BRU™ Month!

In recognition of BRU being chosen "Best Backup Utility" by the readers of *Linux Journal* for the third year in a row, Ted Cook, the President of EST has declared FeBRUary as international BRU month.

When asked to further explain this decree, Mr. Cook stated "Well, with the Readers' Choice Award announcement and the fact that the name 'February' already contained the word 'BRU,' it was a simple decision to make. In fact, it seems that even the inventors of the calendar had pre-ordained that this declaration should be made. I invite all Linux users to examine their current backup strategies and ensure that they have verifiable backups that they can actually recover with."

EST has provided BRU for the Linux world since the operating system was still in its infancy. Since July of 1994, Linux users have had a truly reliable and easy to use backup and restore application available to safeguard their data.

With this, the third Readers' Choice Award for BRU as the "Best Backup Utility," the Linux user community has once again proven that reliable backup is an essential part of a Linux user's administration toolkit.



Thank You For Making BRU #1 ... Again!

In Celebration of Official BRU month, any new purchases of the full release of BRU for Linux made prior to February 28th get a 25% discount off of the normal price of \$299.

Order BRU for Linux Today and Save 25%!

Order Must Be Placed Between Jan 15th and February 28th, 1998. New Orders Only.

 **Enhanced Software
Technologies, Inc.**
Software You Can Trust™




Enhanced Software Technologies, Inc.
"The BRU Guys"

4014 East Broadway Rd., Suite 405 Phoenix, AZ 85040
Sales: (800) 998-8649 Tel: (602) 470-1115 Fax: (602) 470-1116
Email: ljinfo@estinc.com Web: www.estinc.com

> FEATURES

10 University of Toronto WearComp Linux Project




A look at wearable computers and personal privacy issues.
Dr. Steve Mann

20 Hunting Hurricanes

A story of the Scanning Radar Altimeter and Linux.

C. Wayne Wright and Edward J. Walsh



32 COAS: A Flexible Approach to System Administration Tools

A new easy-to-use configuration tool.

Olaf Kirch

38 Csound for Linux

A look at what's happening in the Linux sound world.

David Phillips

> COLUMNS

52 Focus on Software

David A. Bandel

54 System Administration: Caching the Web, Part 2

David Guerrero

82 At the Forge: Creating a Web-based BBS, Part 2

Reuven M. Lerner

93 Linux Apprentice: The login Process

Andy Vaught

> NEWS & ARTICLES

60 KDE: The Highway Ahead

Kalle Dalheimer



64 The GNOME Project

Miguel de Icaza

72 Virtual Network Computing

Brian Harvey



76 Configuring ATM Networks

Wayne J. Salamon

> REVIEW

50 P-Synch: Changing the Way We Change Passwords

Tim Parker

> DEPARTMENTS

6 Letters to the Editor

8 Stop the Presses: Announcements by Sun Troll Tech

Marjorie Richardson

46 Guest Editorial: Software Libre and Commercial Viability

Alessandro Rubini

81 Advertisers Index

88 New Products



> DEPARTMENTS

90 Best of Technical Support

95 Marketplace

> STRICTLY ON-LINE, <http://www.linuxjournal.com/issue58/>

Color Reactiveness on the Desktop

Bowie Poag

U Interviews Informix's Janet Smith

Marjorie Richardson

Book Review: Building Network Management Tools with Tcl/Tk

Syd Logan

> UPCOMING

Internationalization and Localization

Japanization of Linux

The ARGO Autonomous Vehicle

Mediated Reality

> WHAT IS LINUX?

Linux is a multi-user, multi-tasking operating system that runs on many platforms, including Intel processors, Motorola MC68K and DEC Alphas. It implements a superset of the POSIX standard. Linux interoperates well with other operating systems, including Apple, Microsoft and Novell.

The Linux operating system is freely available—it can be copied and redistributed without fees or royalties. The source code for Linux is available on the Internet to anyone who wants it.

For additional information, see <http://www.linuxresources.com/what.html>.

LINUXTM JOURNAL

Managing Editor:

Marjorie Richardson, ljeditor@ssc.com

Editorial Assistants:

Ellen Dahl, Darcy Whitman

Feature Writers:

Olaf Kirch, Dr. Steve Mann, David Phillips,
Edward J. Walsh, C. Wayne Wright

Copy Editors:

David A. Bandel, Dwight Johnson, Jason
Kroll, Michael Shappe, Laurie Tucker,
David Wright

Proofreader:

Riley P. Richardson

Graphic Design and Layout:

Lydia Kinata, lydia@ssc.com

Imaging and Printing:

Brown Printing Company

SSC

Linux Journal is published by
Specialized Systems Consultants, Inc. (SSC)

Publisher:

Phil Hughes, phil@ssc.com

Linux Journal is a registered trade name of
Specialized Systems Consultants, Inc.
LINUX is a registered trademark of Linus Torvalds.

LINUX JOURNAL (ISSN 1075-3583) is published monthly by SSC, 7723 24th Ave. NW, Seattle, WA 98117. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$22/year in the United States, \$27/year in Canada and Mexico, \$37/year elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, P.O. Box 500, Missouri City, TX 77459-9903. Subscriptions start with the next issue. Back issues, if available, are \$6 each, \$8 airmail.

Subscription Information:

E-mail subs@ssc.com

URL www.linuxjournal.com

Phone (888) 66-LINUX

FAX (281) 261-5999

Mail P.O. Box 500

Missouri City, TX 77459-9903 USA

Please allow 6-8 weeks for processing address changes and orders.

Magazine Reseller Information:

E-mail dist@ssc.com

Advertising Information:

E-mail ads@ssc.com


Phone 206-782-7733

Reprint Information:

E-mail linux@ssc.com

Use of Material Published in Linux Journal

In order to protect everyone involved but still allow for the maximum dissemination of material which appears in *Linux Journal*, we have established the following guidelines for reprinting. The author may use articles immediately if he/she includes the phrase "reprinted with permission of *Linux Journal*". Anyone other than the author should contact *Linux Journal* concerning reuse of the material. Generally permission will be granted unless restricted by the author or prior copyright.



About the Cover:

Our cover this month features a photograph taken by Dr. Steve Mann of his wearable computer gear. See feature article "University of Toronto WearComp Linux Project". Beside wearable camcorders, Dr. Mann is experimenting with sensors sewn into clothing to aid the visually impaired. The future is here—we will be assimilated.

It's a shame they banned

Human Cloning.



All Aspen Systems workstations and servers come with unlimited, toll-free technical support, 30-day money-back guarantee and 2-year parts and labor warranty.

Durango II

CPU: Alpha 21164A up to 600 MHz (upgradable)
SDRAM Memory Subsystem: 128-bit data bus supported
L3 Cache: 2 MB cache size
PCI Bus: Four full-length PCI slots (two 64-bit, two 32-bit)
ISA Bus: Two full-length, dedicated ISA slots

Montrose


CPU: Alpha 21164PC up to 533 MHz
SDRAM Memory Subsystem: 128-bit data bus supported
L3 Cache: 1 MB cache size
PCI Bus: Four full-length PCI slots (two 64-bit, two 32-bit)
ISA Bus: Two full-length, dedicated ISA slots

BYTE BEST



Five-Star Rating for Technology and Performance

Byte Magazine, February 1998
Aspen Systems Durango II



... a shame for our competitors, that is. Because *different* people have *different* needs. Some hardware vendors get their boxes working with one Linux version and call themselves *Linux specialists*. Specialists, that is, until you call wanting the newest kernel or the latest driver – the fast talk and confusion on the other end of the line makes you wonder whether they're really as *special* as advertised. Truth is, most vendors just don't get that "Linux user" is an adjective to describe a person with advanced computing requirements, not a noun to describe *their* ideal customer.

At Aspen Systems, we *are different!* While our systems are as competitively priced as those other guys, we have literally years of experience in building, testing, and debugging Linux systems.

We keep pace with the latest in Linux development and create systems that address your individual applications. Which makes our custom systems true performance boosters that will help you get your work done better and faster.

From rackmounts to servers to multi-boot systems to RAID to custom and OEM motherboards, we'll put together the perfect performance fit at a price that's right.

Aspen Systems Workstations and Servers:
Get UnCloned™

Give us a call or check out our website and discover how we can help you meet your individual needs.

AlphaPowered®

www.aspsys.com

1-800-992-9242

Aspen Systems
UnCloning Linux Since 1994

3900 Youngfield Street • Wheat Ridge, Colorado 80033-3865 USA email aspen@aspsys.com Tel +01(303)431-4606 Fax +01(303)431-7196

© 1998 Aspen Systems, Inc. All trademarks are the property of their respective holders.

Applixware 4.4.1

I recently purchased Applixware 4.4.1. My reason for purchasing it was so that I could interoperate with co-workers who run MS Office. Applixware has the best MS Office import filters I have seen, but they are not quite perfect. I found that while text and tables seem to import well, figures are another story. So far, I have only tried importing a few MS Word 7.0 documents, but in general, the figures fail to convert. I recognize the complexity of the conversion task, and I applaud Applix for doing as well as it does. I hope the filters continue to mature, so that I can finally dump Windows 95 off my machine and use Linux exclusively.

--

Steve Falco
sfalco@worldnet.att.net

Erratum in *LJ* February 1998

I know that it is late to correct an erratum in the February issue; nevertheless, I think *Linux Journal* is a magazine to read and save for later use, so even such a late correction could be of some benefit.

The error is in Listing 5 of the article "Attaching Files to Forms", in the column "At the Forge" on page 93. This listing should contain the following line between lines 2 and 3:

```
no strict "refs";
```

If this line is not present, the Perl interpreter will abort the script as soon as the variable reference `$userfile` is used for write. The uploaded file will be created but not written, i.e., it will have zero length.

--

Aldo Mozzi
medico.red@interbusiness.it

Answer to User's Question

In the November *LJ*, "Best of Technical Support" had a question regarding installing Linux (specifically X) on an IBM Thinkpad 365. There is an excellent article regarding this in *Linux Gazette*, <http://www.ssc.com/lg/issue21/notebook.html>.

I was able to get X working on Thinkpad using this guide, and the author, Sam Trenholme, was extremely helpful in getting me over a few problems when I contacted him via e-mail.

--

Nate Dutra
nate@the-wall.net

Eiffel, Design by Contract

Your fine web site, <http://www.linuxresources.com/>, is usually the first thing I check on weekdays when I boot up. You have helped me greatly in learning about Linux.

Thank you very much for reporting on Dr. Meyer, Eiffel and Design by Contract. Here are three open-source Eiffel and Design by Contract resources:

SmallEiffel, the official GNU Eiffel: <http://www.loria.fr/projets/SmallEiffel/>
TOM, a new GPL/LGPL OO language with multiple inheritance and Design by Contract traits: <http://gerbil.org/tom/>
A page I edit: http://www.newhoo.com/Computers/Programming_Languages/Eiffel/

--

Jerry Fass
fass@pitnet.net

Linux Installation and the Open Source Process

Last week I decided to totally change over to Linux after reading the latest horror story about Windows. Apparently, people connected to the Internet on W-95-98 can be snooped on simply by typing their IP number and a few backslashes. Suddenly, their whole system appears in the other person's MS Explorer window.

So here I am, scurrying to get a faster motherboard and a bigger hard drive. I am all enthusiastic about the new KDE and GNOME projects, yet when I read the installation manuals—S.u.S.E. or Red Hat—I am horrified by how much *totally new* and *alien* system configuration is needed for a new Linux user. The problem is exemplified by the marvelous advertising flyer sent out by S.u.S.E. for its 5.2 release. It seduces the user with a DOS-Windows box packed with a lifetime accumulation of precious utilities. It implies that UNIX clones for these treasures (and more) effortlessly await on the new Linux user's desktop. The actual case is that this happens only after endless installations and configurations.

The open-source concept has been much in the news lately, yet it seems that these installation processes are the one place where the open-source

environment is not used to evolve solutions to these problems. Rather, it is the special province of the private bundling companies. Regardless of whether they put their products in the public domain, they are not developed in the open. People seem to believe that the greed of these companies will produce the fastest results, but I have not seen any miracles yet. Perhaps there are installation projects underway in the open-source community. If so, nothing could be more critical to the advancement of Linux.

We often hear people crowing in *LJ* that some huge corporation or the defense department is now more efficient because of Linux. I am much more impressed by ease of use by the ordinary person. If *LJ* were to make open-source installation projects a continuing focus of articles, it could do an incredible service to the evolution and spread of Linux.

--

David Briars
dbriars@sover.net

Easier installation seems to be what everyone is asking for these days. Red Hat is working on it with LinuxConf, and Caldera with COAS (see article in this issue).

--Editor

Happy Hacking Keyboard Price

The correct price for the Happy Hacking Keyboard at the time of the article was \$159, not \$189. Today I cut the price again from \$159 to \$139 as a standard price. Please let your readers know.

--

Ted Abe, PFU America, Inc.
tetsu@pfuca.com

BTS Correction

In the November 1998 "Best of Technical Support", weird things have been happening to Eric Benoit's system. **su** misbehaves, as do **man** and **less**. Scott Maxwell thinks it may be terminal options, but the answer lies with `/dev/tty`. My bet is that somehow, something has changed the permissions on it and it is no longer readable. Just type:

```
chmod u+rw /dev/tty
```

and all will be well. You should probably make sure it is writable

Letters continued on page 94

Let Microway build your next Linux, UNIX or NT
Workstation, Server or Beowulf Cluster using...

Screamer™ 667


Cray Performance at a PC Price!

667 MHz Alpha, 1.3 Gigafllops, 8MB Cache

**Dual Alpha EV6
2000 Megaflops
performance
now shipping!
Call Microway for
a quotation.**

Since 1982 Microway has provided the PC world with the fastest numeric devices and software available. No product in the last 17 years has excited us more than the Alpha Screamer. With its ability to execute 2.7 billion operations per second, the Screamer is the best choice for your next workstation or server! In addition to Linux and Beowulf applications, the Screamer runs NT, Digital UNIX and Open VMS.

This means you can run PVM and MPI Beowulf applications on the same hardware that runs Microsoft Excel or Word, Oracle, Fortran, C/C++ and Visual Basic. Also supported are Pro/Engineer, MicroStation, ANSYS, LAPACK, Gaussian, Softimage and Lightwave. Over the last 17 years we have designed systems for thousands of satisfied customers worldwide, including major universities, Fidelity Investments, Lucent Technologies and NASA. Our technicians are expert at configuring all Alpha operating systems, and you will not find more technically competent sales people anywhere.



Custom Screamer Workstations

Microway systems include 1 to 4 CPUs with fast caches, up to 4GB of high-speed memory, 6 PCI slots, SCSI hard drives, 3D graphics cards, DLT drives and libraries, and RAID subsystems. Microway's exclusive 8MB SSRAM cache, fed by a 288-bit wide memory system, boosts performance by up to 100%. Screamer™ workstations range in price from \$2,795 to over \$80,000.

Microway also produces one of the finest numeric optimized compilers - NDP Fortran. Since 1986, hundreds of applications have been ported to the x86 and Alpha with it. Using hand-coded BLA's and FFT's, our NDP VDSP Alpha Library hits 560 megaflops triangularizing dense arrays and performs a 1024 complex FFT in 146 microseconds. This library also includes all LAPACK subroutines.

AlphaPowered™

Press Accolades for Microway's Screamer

Windows NT Magazine Lab Report - August, 1998

"On the AIM WNT Peak Performance metric. . . Microway scored 643.7. It scored higher on the AIM WNT Sustained Performance metric-219.4 than any other Alpha NT System I've tested."

LINUX Journal - January, 1998

"Literally everything runs profoundly faster on the Screamer."

Windows Sources - February, 1998

"the Microway system blew away the best Intel-based workstations we've tested. . . on our number-crunching Lightwave 3D test."

PC Computing - July, 1997 ★★★★★

Microway's Screamer . . . "is, quite simply, the fastest Windows NT machine on the planet. . . The performance leader."

Computers in Physics - September, 1997 Product of the Month

Visit **www.microway.com** for complete product information or call technical sales at **508-746-7341**.

Digital, Alpha, and Digital UNIX TM Digital.
Visual Basic, NT, Excel and Word TM Microsoft.
Screamer, NDP Fortran and Microway TM Microway.

Microway®

Technology You Can Count On

Corporate Headquarters: Research Park, Box 79, Kingston, MA 02364 USA • TEL 508-746-7341 • FAX 508-746-4678 • www.microway.com
info@microway.com • Australia 61 292094580 • Denmark 45 39624156 • Germany 49 6976752384 • India 91 806637770 • Italy 39 290782776
Japan 81 64593113 • Korea 82 25561257 • New Zealand 64 33595556 • Poland 48 22487172 • Spain 34 35809444 • UAE 97 19281081

Announcements by Sun and Troll Tech

by Marjorie Richardson

On December 8, Sun Microsystems made two announcements of interest to the Linux community. One was the completion of the Linux port to the UltraSPARC architecture; the other was the new, more open licensing of Java.

UltraSPARC

When Sun joined Linux International back in May, it was with the expressed intention of joining the Linux community to do the UltraSPARC port. This has now become a reality. In addition, they have announced their intention to allow vendors to sell the UltraSPARC preloaded with Linux as well as Solaris.

Every machine sold preloaded with Linux is another win for Linux. An even bigger win is having yet one more of the "big guys" acknowledge that computers with Linux pre-installed are more attractive to potential buyers, especially those new to Linux. I for one am happy to see Sun following in the footsteps of Corel Computer and Cobalt Networks in making this decision.

Java Licensing

The new, open licensing for Java has been speculated about for some time. Will Sun make it open? If so, when? Well, they did it with this announcement—another big win for the Open Source movement. Source code has always been free for non-commercial use and the binaries have been freely available for use in tools developed by others. Here's how it has become more open, according to the press release:

- Allows commercial entities to use and modify the source code for commercial software product development without charge.

- Allows innovation of the source code without requiring that innovation be returned to Sun.
- Allows commercial entities to modify and share compatible source code with other commercial entities without charge and without mediation from Sun.
- Allows licensees to package for resale Sun's Java platform class libraries with virtual machines from other licensees.

These are major changes, but not quite the GPL. Developers who actually incorporate the code into a commercial product will still be required to pay a fee to Sun. Still, it's a step in the right direction and others are sure to follow suit.


Troll Tech

In a similar vein, Troll Tech announced in November that it plans to release version 2.0 of the Free Edition of the Qt graphical interface under an Open Source license. This will eliminate any worries and controversy regarding the inclusion of KDE desktop in commercial products.

Good news, indeed, to everyone who has wished for a user-friendly desktop for Linux. KDE has come a long way toward providing that option for those who are shy of the command line. (See "KDE: The Highway Ahead" by Kalle Dalheimer in this issue.)

In our ongoing quest to bring enjoyment to our readers, we introduce the Linux-based comic strip *User Friendly* by Illiad: <http://www.userfriendly.org/>.

USER FRIENDLY by Illiad



**Like Mom always told you,
teamwork will get you further
in life...**

**Now it will get you the best preinstalled
Linux solution on Earth!**

**S.u.S.E. and VA Research combine
the best of the Linux world**




Since 1993, S.u.S.E. has provided Europe with cutting edge Linux software solutions like S.u.S.E. Linux, long Europe's favorite distribution. Known for advanced German-engineered tools like the YaST system administration tool, S.u.S.E. Linux exceeds the highest of expectations.

For just as long, VA Research, based in Silicon Valley, has been the name synonymous with premier Linux hardware solutions. Whether it's servers, workstations or laptops, VA Research always delivers top performance, reliability, value and service.

The result is a better solution for you

At S.u.S.E. and VA Research, we've combined the best of our expertise, which results in an even better Linux solution. As you certainly know, especially when you work with newer, advanced hardware, optimizing your Linux system can be a challenge. Our mission is to eliminate those hassles and surprises and deliver top performance right out of the box. We can deliver because we've done our homework, benchmarked our hardware, and realized maximum compatibility. Furthermore, we provide solutions like the XSuSE X Servers, which support an ever-broader range of high-end graphics cards.



Let our team help your team

Whether your needs call for Linux solutions for home or office, put the S.u.S.E./VA Research Team to work for you. Our knowledgeable staffs are standing by, ready to assist you.



For S.u.S.E. Linux inquiries:

S.u.S.E., Inc.

Toll-free 1-888-875-4689

1-510-835-7873

Fax +1-510-835-7875

<http://www.suse.com>

<http://www.suse.de> (German)

info@suse.com

suse@suse.de (in Europe)



For preinstalled S.u.S.E. Linux systems:

VA Research

Toll-free 1.888.LINUX4U

1-650-934-3666

Fax +1-650-964-7668

<http://www.varesearch.com>

sales@varesearch.com

University of Toronto WearComp Linux Project

Dr. Mann describes his WearComp ("Wearable Computer") invention and how it has evolved into the same kind of philosophical basis for self determination and mastery over one's own destiny that is characteristic of the Linux operating system currently running on WearComp.

by Dr. Steve Mann

This paper is part one of a two-part series. In this part I will describe a framework for machine intelligence that arises from the existence of human intelligence in the feedback loop of a computational process.

I will also describe the apparatus of the invention that realizes this form of intelligence, beginning with a historical perspective outlining its visual and photographic origins. The apparatus of this invention, called "WearComp", emphasizes self-determination and personal empowerment.

I also intend to present the material within a philosophical context I call COSHER (Completely Open Source, Headers, Engineering and Research) that also emphasizes self-determination and mastery over one's own destiny.

This "personal empowerment" aspect of my work is what I believe to be a fundamental issue in operating systems such as Linux. It is this aspect that WearComp and Linux have in common, and it is for this reason that Linux is the selected operating system for WearComp.

An important goal of being COSHER is allowing anyone the option of acquiring, and thus advancing, the world's knowledge base.

I will also introduce a construct called "Humanistic Intelligence" (HI). HI is motivated by the philosophy of science, e.g., open peer review and the ability to construct one's own experimental space. HI provides a new synergy between humans and machines that seeks to involve the human rather than having computers emulate human thought or replace humans. Particular goals of HI are human involvement at the individual level and providing individuals with tools to challenge society's preconceived notions of human-computer relationships. An emphasis in this article is on computational frameworks surrounding "visual intelligence" devices, such as video cameras interfaced to computer systems.

Problem Statement

I begin with a statement of what I believe to be a fundamental problem we face in today's society as it pertains to computers and, in particular, to computer program source code and disclosure. Later, I will suggest what I believe to be solutions to this problem. Linux is one solution, together with an outlook based on science and on self-determination and individual empowerment at the personal level.

A first, fundamental problem is that of software hegemony, seamlessness of thought and the building of computer science upon a foundation of secrecy. Advanced computer systems is an area where a single individual can make a tremendous contribution to the advancement of human




Figure 1. Evolution of the WearComp Invention

knowledge, but is often prevented from doing so by various forms of software fascism. A system that excludes any individual from exploring it fully may prevent that individual from "thinking outside the box" (especially when the box is "welded shut"). Such software hegemonies can prevent some individuals from participating in the culture of computer science and the advancement of the state of the art.

A second fundamental problem pertains to some of the new directions in human-computer interaction (HCI). These new directions are characterized by computers everywhere, constantly monitoring our activities and responding *intelligently*. This is the ubiquitous surveillance paradigm in which keyboards and mice are replaced by cameras and microphones watching us at all times. Perpetrators of this environmental intelligence claim we are being watched for our benefit and that they are making the world a better place for us.

Computers everywhere constantly monitoring our activities and responding intelligently have the potential to make matters worse from the software hegemony perspective, because of the possibility of excluding the individual user from knowledge not only of certain aspects of the computer upon his or her desk, but also of the principle of operation and the function of everyday things. Moreover, the implications of secrecy within the context of these intelligence-gathering functions puts forth a serious threat to personal privacy, solitude and freedom.



Computer Science or Computer Secrecy

Science provides us with ever-changing schools of thought, opinions, ideas and the like, while building upon a foundation of verifiable (and sometimes evolving) truth. The foundations, laws and theories of science, although

Dear Cisco® 2501 users, try the Cyclades®-PR3000.

Get performance and flexibility,

The **Cyclades-PR3000** is a powerful, midrange router. The PR3000 offers three configurable WAN interfaces for quick expansion and flexibility.



For a more detailed description of the results presented in these two graphs go to our web site at:
http://www.cyclades.com/pr3k_ads1.html

See what others have to say,

“As Colorado’s largest ISP we purchase for our own use and recommend Cyclades to our customers, because of its performance, reliability, and the strength of their technical support.”

W. R. Sheppard, CEO
BeWell.Net
June 3rd, 1998

Visit our website:
www.cyclades.com




Cyclades-PR3000-TS with 32 ports.

Cyclades-PR3000

- Midrange, Expandable Router
- IP, IPX, SNA
- Terminal Server Support

Choice of cards to install in the PR3000:



- T1/E1 Interface Card with on board CSU/DSU
- ISDN BRI Interface Card
- Serial Interface Card
- Terminal Server Card (Z-Bus Interface) for up to 64 serial ports.



Fremont, CA USA
Tel: 510-770-9727
1-800-882-9252
FAX: 510-770-0355
sales@cyclades.com



true by assumption, may at any time be called into question as new experimental results unfold. Thus, when doing an experiment, we may begin by making certain assumptions; at any time, these assumptions may be verified.

In particular, a scientific experiment is a form of investigation that leads wherever the evidence may take us. In many cases, the evidence takes us back to questioning the very assumptions and foundations we had previously taken as truth. In some cases, instead of making a new discovery along the lines anticipated by previous scientists, we learn that another previous discovery was false or inaccurate. Sometimes these are the biggest and most important discoveries—things that are found out by accident.

Any scientific system that tries to anticipate “what 99% of the users of our result will need” may be constructing a thought prison for the other 1% of users who are the very people most likely to advance human knowledge. In many ways, the entire user base is in this thought prison, but many would never know it since their own explorations do not take them to the outermost walls of this thought prison.

Thus, a situation in which one or more of the foundation elements are held in secret is contrary to the principles of science. Although many results in science are treated as a “black box”, for operational simplicity there is always the possibility that the evidence may want to lead us inside that box.

Imagine, for example, conducting an experiment on a chemical reaction between a proprietary solution “A”, mixed with a secret powder “B”, brought to a temperature of 212 degrees T. (Top-secret temperature scale which you are not allowed to convert to other units.) It is hard to imagine where one might publish results of such an experiment, except perhaps in the *Journal of Non-Reproducible Results*.

Now, it is quite likely that one could make some new discoveries about the chemical reaction between A and B without knowing what A and B are. One might even be able to complete a doctoral dissertation and obtain a Ph.D. for the study of the reaction between A and B (assuming large enough quantities of A and B were available).

Results in computer science that are based, in part, on undisclosed matters inhibit the ability of the scientist to follow the evidence wherever it may lead. Even in a situation where the evidence does not lead inside one of the secret “black boxes”, science conducted in this manner is irresponsible in the sense that another scientist in the future may wish to build upon the result and may, in fact, conduct an experiment that leads backwards as well as forwards. Should the new scientist follow evidence that leads backwards, inside one of these secret black boxes, then the first scientist will have created a foundation contaminated by secrecy. In the interest of academic integrity, better sci-

ence would result if all the foundations upon which it was built were subject to full examination by any scientist who might, at some time in the future, wish to build upon a given discovery.

Thus, although many computer scientists may work at a high level, there would be great merit in a computational foundation open to examination by others, even if the particular scientist using the computational foundation does not wish to examine it. For example, the designer of a high-level numerical algorithm who uses a computer with a fully disclosed operating system (such as Linux) does other scientists a great service, even if he uses it only at the API level and never intends to look at its source code or that of the Linux operating system underneath it.

Obvious or Obfuscated

Imagine a clock designed so that when the cover was lifted off, all the gears would fly out in different directions, such that a young child could not open up his or her parents’ clock and determine how it works. Devices made in this manner would not be good for society, in particular for the growth and development of young engineers and scientists with a natural curiosity about the world around them.

As the boundary between software and hardware blurs, devices are becoming more and more difficult to understand. This difficulty arises in part as a result of deliberate obfuscation by product manufacturers. More and more devices contain general-purpose microprocessors, so that their function depends on software. Specificity of function is achieved through specificity of software rather than specificity of physical form. By manufacturing everyday devices in which only executable code is provided, manufacturers have provided a first level of obfuscation. Furthermore, additional obfuscation tools are often used in order to make the executable task image more difficult to understand. These tools include

strippers that remove things such as object link names and even tools for building encrypted executables which contain a dynamic decryption function that generates a narrow sliding window of unencrypted executable, so that only a small fragment of the executable is decrypted at any given time. In this way, not only is the end user deprived of source code, but the executable code itself is encrypted, making it difficult or impossible to look at the code even at the machine-code level.

Moreover, complex programmable logic devices (CPLDs), such as the Alterra 7000 series, often have provisions to permanently destroy the data and address lines leading into a device, so that a single chip device can operate as a finite-state machine yet conceal even its machine-level contents from examination. (See Resources 1 for an excellent tutorial on FPGAs and CPLDs.) Devices such as




Figure 2. ECE1766 Class Pictures

Cyclades Multiport Cards make it simple.

High speeds, easy to configure, and cost effective. *Why make it hard on yourself?*

Remote Access (RAS) Home Automation Test Bed Systems Industrial Automation ISP Applications

"Just plug in some more serial ports," says your Boss. "Nothing simpler."

Usually life isn't that simple. . . till now.

Cyclades Multiport Cards are simple to install and easy to maintain for your applications. Just plug the host card and in minutes you will have more serial ports to connect to. From creating a Remote Access Server to Home Automation you will not find anything easier to setup than a Cyclades Multiport Card.

Adding more serial ports will either be a "piece of cake", or the beginning of a really long week. Make it a short week by using Cyclades Multiport Cards in your next upgrade.

High-end Solution


Cyclades-Ze Series [16-64 ports 920Kbps]
Intelligent, Rack Mountable, Near zero CPU overhead

Entry Level Solution

Cyclom-Y Series [8-32 ports, 115Kbps]
RISC-based, Cost-effective solution



Cyclades-Ze
Rack Mountable & Expandable



Cyclom-8Yo
8 Port w/ Octopus cable

All Cyclades Multiport Cards are backed by a Lifetime Warranty from manufacturing defects and free driver updates via the Internet.

Satisfied Customers of Cyclades:




MOTOROLA



PAGENET®

Polaroid

Drivers for Windows NT, Windows 95, Windows 98, Linux, FreeBSD, BSDI, SCO
PCI and ISA, RJ-45, DB-25, DB-9, Surge Protection



www.cyclades.com

Fremont, CA USA Tel: 510-770-9727 1-800-882-9252 FAX: 510-770-0355
sales@cyclades.com



Clipper chips go a step further by incorporating fluorine atoms, so that if the user attempts to put the device into a milling machine to mill it off layer by layer for examination under an electron microscope, the device will self-destruct in a quite drastic manner. Thus, the Clipper phones could contain a “Trojan horse” or some other kind of back door and we might never be able to determine whether or not this is the case—yet another example of deliberate obfuscation of the operational principles of everyday things.

We have a growing number of general-purpose devices in which the function or purpose depends on software, downloaded code or microcode. Because this code is intellectually encrypted, so is the purpose and function of the device. In this way, manufacturers may provide us with a stated function or purpose, but the actual function or purpose may differ or include extra features of which we are not aware.

Environmental Intelligence Gathering Systems

A number of researchers have been proposing new computer user interfaces based on environmental sensors. Buxton, who did much of the early pioneering research into intelligent environments (smart rooms, etc.), was inspired by automatic flush urinals (as described, for example, in U.S. Pat. 4309781, 5170514, etc.) and formulated, designed and built a human-computer interaction system called the “Reactive Room” (see Resources 2 and 3). This system consisted of various sensors, including optical sensors (such as video cameras) and processing, so that the room would respond to the user’s movement and activity.


Increasingly, we are witnessing the emergence of intelligent highways, smart rooms, smart floors, smart ceilings, smart toilets, smart elevators, smart light switches, etc. However, a typical attribute of these “smart spaces” is that they were designed by someone other than the occupant. Thus, the end user of the space often does not have a full disclosure of the operational characteristics of the sensory apparatus and the flow of intelligence data from the sensory apparatus.

In addition to the intellectual encryption described in the previous section, where manufacturers could make it difficult, or perhaps impossible, for the end user to disassemble such sensory units in order to determine their actual function. There is also the growth of hidden intelligence, in which the user may not even be aware of the sensory apparatus. For example, U.S. Pat. 4309781 (for a urinal flushing device) describes:

... sensor... hidden from view and thus discourage tampering with the sensor... when the body moves away from the viewing area... located such that an adult user of average height will not see it... sensing means, will be behind other components... positioned below the solenoid to allow light in and out. But the solenoid acts in the nature of a hood or canopy to shield the sensing means from the normal line of sight of most users.... Thus most users will not be aware of the sensing means. This will aid in discouraging tampering with the sensing means. A possible alternate arrangement would be to place the sensing means below and behind the inlet pipe.

U.S. Pat. 4998673 describes a viewing window concealed inside the nozzle of a shower head, where a fiber

optics system is disclosed as a means of making the sensor remote. The concealment is to prevent users from being aware of its presence. U.S. Pat. 5199639 describes a more advanced system where the beam pattern of the nozzle is adapted to one or more characteristics of the user, while U.S. Pat. 3576277 discloses a similar system based on an array of sensing elements.



A method of creating viewing windows to observe the occupants of a space while at the same time making it difficult for the occupants to know if and when they are observed is proposed in U.S. Pat. 4225881 and U.S. Pat. 5726706.

In addition to concealing the sensory apparatus, a goal of many visual observation systems is to serve the needs of the system architect rather than the occupants. For example, U.S. Pat. 5202666 discloses a system for monitoring employees within a restroom environment, in order to enforce hygiene (washing of hands after using the toilet).

Other forms of intelligence, such as intelligent highways, often have additional unfortunate uses beyond those purported by the installers of the systems. For example, traffic-monitoring cameras were used to round up, detain and execute peaceful protesters in China’s Tiananmen Square.

U.S. Pat. 4614968 discloses a system where a video camera is used to detect smoke by virtue of the fact that smoke reduces the contrast of a fixed pattern opposite the video camera. However, the patent notes that the camera can also be used for other functions such as visual surveillance of an area, since only one segment or line of the camera is needed for smoke detection. Again, the camera may thus be justified for one use; additional uses, not disclosed to occupants of the space, may then evolve. U.S. Pat. 5061977 and 4924416 disclose the use of video cameras to monitor crowds and automatically control lighting in response to the absorption of light by the crowds. While this form of environmental intelligence is purportedly for the benefit of the occupants (to provide them with improved lighting), there are obvious other uses.

U.S. Pat. 5387768 discloses the use of visual inspection of users in and around an automated elevator. Again, these provide simple examples of environmental intelligence in which there are other uses, such as security and surveillance. Although even those other uses (security and surveillance) are purportedly for the benefits of the occupants, and it is often even argued that concealing operational aspects of the system from the occupants is also for their benefit, it is an object of this paper to challenge these assumptions and provide an alternate form of intelligence.

When the operational characteristics, function, data flow and even the very existence of sensory apparatus is concealed from the end user, such as behind the grille of a smoke detector, environmental intelligence does not necessarily represent the best form of human-machine

APPGEN[®]

PowerWindows[™] Business Applications

PowerWindows[™] is a set of fully integrated, modifiable, portable, scalable, consistent, multi-user, multi-company, industrial-strength business applications used and developed worldwide since 1979. They run simultaneously on GUI client PCs and character-based terminals from the same server, with powerful features for the most demanding businesses and organizations.

PowerWindows applications are shipped with source code. They are thoroughly documented, and context-sensitive help is guaranteed always available.

Accounts Receivable
General Ledger
Sales Order Processing
Payroll Processing
Bill of Materials
Bank Reconciliation

Accounts Payable
Inventory Control
Purchase Order Processing
Job Cost Tracking
Professional Billing
and more on the way...

Appgen 4GL

Fast, powerful fourth generation development environment that *SCO World* described as "Installation Ecstasy" is perfect for the development of vertical applications—especially since the basic business apps are all ready and waiting for you.

Beyond 2000
since 1979

Business Apps
with Source Code


Runs on Linux
SCO Unix[®] AIX[®] Solaris[®] Windows NT[®]
and just about everything else
that matters

It's a sweet deal
for Resellers

For info/demo, toll-free
1-877-859-7500

PowerWindows[™] from Appgen Business Software, Inc.

Appgen[®] is a registered trademark and PowerWindows[™] is a trademark of Appgen Business Software, Inc., Bohemia, New York 11716. Other trademarks are the property of their respective owners.



Sold **only** through Appgen authorized
Value-Added Resellers

relationship for all concerned. Even when the sensors are visible, there must be the constant question as to whether or not the interests of the occupant are identical to those who control the intelligence-gathering infrastructure.

The need for personal space, free from monitoring, has also been recognized (see Resources 4) as essential to a healthy life. As more and more personal space is stolen from us, we may need to be the architects of alternate spaces of our own.

Solution to Software Fascism

The first solution to these problems is a framework called Completely Open Source, Headers, Engineering, and Research (COSHER). Before investing considerable time in learning how to use new software and in developing works for that new software, which may then become *locked* into a particular file format, we ask ourselves a very simple question: is the software in question COSHER?

This means that there has been no deliberate attempt at obfuscation of the underlying principles of the operation of this software or in preventing us from freely distributing the intellectual foundations upon which we may invest many years of our lives. Deliberate attempts at obfuscation include such practices as eliminating source code and stripping executable task images.

By using COSHER software, we are making a statement that we prefer Computer Science to Computer Secrecy. Science supports the basic principles of peer review, a continued development and advancement of software principles and principles that we build on top of the software.

Moreover, the time we invest in learning the software as well as creating works in the software will be less likely to go to waste if we have a copy of the complete source code of the software. In this manner, should the software ever become discontinued or unsupported, we will be able to become our own software support group and migrate the software forward to new architectures as our old computers become obsolete. If it is COSHER, chances are we will be less likely to lose the many hours or years we invest in producing works within the software. Furthermore, if we make new discoveries that are built on a foundation of COSHER software, they are easier to distribute.

In science, it is important that others be able to reproduce our results. Imagine what it would be like if we had built our results on top of DOS 3.1. Others would have to either rewrite our software to exactly reproduce our results, or find an old version of DOS 3.1. Since this is proprietary software, we are not at liberty to freely distribute it with our research, but it is also no longer available for purchase. However, if we had built our work on COSHER software such as Linux 1.13, we can include a full distribution of Linux 1.13 in an archive together with our results. Many years in the future, a scientist wishing to reproduce our results could then obtain a virtual machine (emulator for our specific architecture which will no doubt be obsolete by then) and install the COSHER operating system (Linux 1.13) that came with our archive, then compile and run our programs.

The Linux operating system is a good example of a COSHER operating system. GNU software is also COSH-

ER. Many COSHER software packages are available, including GIMP (Gnu Image Manipulation Program) and the VideoOrbits software package (described in <http://wearcam.org/orbits/index.html>).

Solution to Environmental Intelligence Gathering

I propose a computational framework for individual personal empowerment. This framework is based on my "WearComp" invention—an apparatus for (embodiment of) realization of HI.

This framework involves designing a new kind of personal space. An embodiment of the "WearComp" invention is an apparatus that is owned, operated and controlled by the occupant of that space. In one sense, the apparatus of this invention is like a building built for one occupant and collapsed down around that one occupant.

WearComp as a Basis for HI

I invented WearComp in Canada in the 1970s as a photographic tool for the visual arts (see Resources 5), in particular, something I called "mediated reality" (altered perception of visual reality). The goal of mediated reality, unlike related concepts such as virtual (or augmented) reality, was to reconfigure (augment, deliberately diminish or otherwise alter) the perception of reality in order to attain a heightened awareness of how ordinary, everyday objects respond to light.

HI is a new form of human-computer interaction comprising a computer that is subsumed into the personal space of the user (e.g., the computer may be worn, hence the term "user" and "wearer" of the computer are interchangeable), controlled by the wearer, with both operational and interactional constancy (e.g., it is always on and always ready and accessible [see Resources 6]).


The WearComp invention, described in IEEE Computer, Vol. 30, No. 2 at <http://wearcomp.org/ieeecomputer.htm> (a historical account was given in IEEE ISWC-97, October 1997 and is also on-line at <http://wearcomp.org/historical/index.html>) forms the basis for HI. The evolution of the apparatus of this invention is depicted in Figure 1.

Definition of WearComp

A wearable computer is a computer that is subsumed into the personal space of the user, controlled by the user and has both operational and interactional constancy.

Most notably, it is a device that is always with the user and into which the user can always enter commands and execute a set of entered commands while walking around or doing other activities.


The most salient aspect of computers in general (whether wearable or not) is their *reconfigurability* and their *generality*, e.g., their function can be made to vary widely, depending on the instructions provided for program execution. This is true for the wearable computer (WearComp). For example, the wearable computer is more than just a wristwatch or regular eyeglasses; it has the full functionality of a computer system and, in addition, is inextricably intertwined with the wearer.



Announcing the Fastest, Most Affordable RAID Disk Arrays for Linux

New RAIDZONE Technology Gives Unbeatable Performance Using High Capacity Ultra ATA Disks

- Supports mirroring, striping, RAID 5 & RAID 10
- Full hot swap and hot spare capability, including the ability to reconfigure the disk array on the fly, without a reboot
- Exceptional Linux performance as a file server (including Samba), and as a database server



- DirectPCI Technology – every drive has its own disk controller IC which is directly on the 132 MB/Sec PCI Bus
- Full support for Linux 2.2, including an easy-to-use Java-based administration program that automates disk array configuration network-wide
- Uses the latest high speed, high capacity Ultra ATA disks
- Internal, external and rack mount arrays support from 5 to 75 or more disks

RAIDZONE Disk Arrays Done Right

The RAIDZONE family of disk array solutions has been designed from the ground up to provide unmatched performance at unbeatable prices. RAIDZONE's DirectPCI technology completely eliminates the need for a peripheral bus (i.e. SCSI) between disk array and computer.

Every disk in a RAIDZONE array is independently connected directly to the computer's 132 MB/Sec PCI Bus. Complete parallelism enables RAIDZONE arrays to achieve unheard of performance using high capacity Ultra ATA disks at a fraction of the cost of other disk arrays.


All RAIDZONE arrays provide the features and capabilities of disk arrays costing \$30,000 or more. Hot swap, hot spare, plus the ability to reconfigure the disk array on the fly, without a reboot. Install and boot Linux 2.2 directly from the array, or add an array to an existing system. Dual hot-swappable fans are standard. RAIDZONE monitors voltage & temperature levels and disk status, and provides automatic exception notification.

From 32 GB up to 1 Terabyte or more, there is a RAIDZONE solution for all your disk array requirements. Nobody buys just one!

For more information
about RAIDZONE,
visit our website at
www.raidzone.com
or give us a call at
1-800-388-1896
OEM, VAR & Dealer
inquiries welcome.

Have It Your Way: Internal, External, Rack


Can of RAID



- Internal unit holds 5 disks, uses only 3 bays
- Hot swap drives & fans
- Supports multiple cans per computer

Box of RAID

- External unit, 5 disks, single cable to host computer
- 10- and 15-disk units and hot swappable power supplies available



Rack of RAID

- 10-, 15-, and 24-disk racks with hot-swappable power supplies
- Supports multiple racks per computer



Stack of RAID



- Up to 60 disks, 4 rack mount enclosures per computer
- Configurations of up to a Terabyte or more, high performance, fully redundant

RAIDZONE and DirectPCI are trademarks of Consensus Corp.

This is what sets the wearable computer apart from other wearable devices such as wristwatches, regular eyeglasses, wearable radios, etc. Unlike these other wearable devices that are not programmable (reconfigurable), the wearable computer is as reconfigurable as the familiar desktop or mainframe computer.

The formal definition of wearable computing defined in terms of its three basic modes of operation and its six fundamental attributes is provided elsewhere in the literature. (See Resources 7.)

WearComp, as Universal Interface to Reality

Such a computational framework allows one to subsume all of the personal electronics devices one might normally carry, such as cellular phone, pager, wrist watch, heart monitor, camera and video camera into a single device. Obviously, since it is a fully featured computer, it is possible to respond to e-mail, plan events on a calendar, type a report, etc., while walking, standing in line at the bank or anywhere. In this way, WearComp anticipated the later arrival of the so-called "laptop computer", but has advantages over the laptop in the sense that it can be used while walking around doing other things. However, the real power of WearComp is in its ability to serve as a basis for personal imaging and humanistic intelligence.

Personal Safety Device

WearComp not only subsumes the function of the laptop computer, but goes beyond it. Another area in which WearComp provides a truly new form of user interface not found on laptops and PDAs (personal digital assistants) is in its constancy of user interface and operation. This characteristic may become most evident in its use as a personal security camera. Imagine, perhaps as you walk down some quiet street at night, an assailant appears, demanding cash from you. You would not likely have the time or opportunity to pull out a camcorder to record the experience, but since the eyeglasses are worn constantly, you would have a video record of the experience to aid investigation.

Camera of the Future

Less extreme examples of WearComp as a new user-interface include the ability to construct a personal documentary video without conscious thought or effort. For example, in a fully mediated reality, all light entering the eyes, in effect, passes through the computer and may therefore be recorded (and possibly transmitted to remote locations). Wearable Wireless Webcam (see Resources 8) is an example of a personal documentary video recorded using a reality mediator.

In the future, we may very well have the capability to capture and recall our own personal experiences and to have photo albums generated automatically for us. We will never miss baby's first steps, because we will have a retroactive record feature that lets us, for example, "begin record-

ing from 5 minutes ago". Photo albums, in addition to being generated automatically, may also be exhibited while they are being generated. Rather than sending postcards to friends and relatives or showing them an album after you come back from vacation, you may just put on your sunglasses and have the album sent to them automatically, as was done with the Wearable Wireless Webcam experiment in which video was transmitted and still images automatically selected from the video.

Personal Intelligence Arms Race

While there will no doubt be more environmental intelligence than personal intelligence, there is at least the hope that there might be an end to the drastic imbalance between the two. The individual making a purchase in a department store may have several cameras pointing at him to make sure that if he removed merchandise without payment, there

would be evidence of the theft.

However, in the future, he will have a means of collecting evidence that he did pay for the item, or a recorded statement from a clerk about the refund policy. More extreme examples such as the case of Latasha Harlins, a customer falsely accused of shoplifting and fatally shot in the back by a shopkeeper as she attempted to walk out of the shop, come to mind.

In this sense, the camera-based reality mediator becomes an equalizer much like the Colt 45 in the "Wild West". In the WearCam case, it is simply a matter of mutually assured accountability.

Future Directions


Much work remains to be done in development of this project. Currently, I teach Electrical and Computer Engineering (ECE1766) at the University of Toronto. To

the best of my knowledge, this is the world's first course on how to be a "cyborg" entity. Students learn not only by doing, but by *being*. I call this form of learning *existential learning*. Each student creates a "reconfigured self"—a new form of personal space. Thus, students learn about the concept of personal empowerment from a first-person perspective through personal involvement.

We are writing new protocols for the altered perception of reality (mediated reality) that the WearComp provides. One example is picture-transfer protocol (PTP), in which packets of variable length are transmitted. Each packet is a JPEG compressed picture. Because of image compression, the amount of data varies depending on image content, hence the packet length depends on image content.

The reason for one packet per picture is that pictures are taken 60 times per second, which is much faster than they can be sent. Thus, whenever there is a lost packet and a re-transmission is needed, a newer picture will most likely be available to be sent instead. With PTP, retransmissions are always current.

Next month I will describe a mathematical (computational) framework called "Mediated Reality", in which we will see that picture data is of greatest value only if it is



up-to-date. Old pictures are of less value when trying to construct a computer-mediated reality. Thus, packet resends should always be of the most current image; hence the design of PTP is based on variable packet lengths, in which the packet length is the length of a picture.

Further information about the WearComp Linux project may be found in <http://wearcam.org/ece1766.html>. ■

Steve Mann, inventor of WearComp (wearable computer) and WearCam (eye-tap camera and reality mediator), is a faculty member at the University of Toronto, Department of Electrical and Computer Engineering. Dr. Mann has been working on his WearComp invention for more than 20 years, dating back to his high school days in the 1970s. He brought his inventions and ideas to the Massachusetts Institute of Technology in 1991, founding what later became the MIT Wearable Computing Project, and received his Ph.D. from MIT in 1997 in this new field he had established. Anyone interested in joining or helping out with the "community of cyborgs" project or the WearComp Linux project may contact the author by e-mail at mann@eecg.toronto.edu.

Thanks to Kodak and Digital Equipment Corporation (DEC) for assistance with the Personal Imaging and Humanistic Intelligence projects.

Resources

1. *IEEE Design and Test of Computers*, 1996, "Fpga and cpld architectures: A tutorial", Stephen Brown and Jonathan Rose. See also <http://wearcam.org/jayarpubs.html>.
2. "Reactive room", 1995, Jeremy Cooperstock and Bill Buxton, <http://www.dgp.utoronto.ca:80/people/rroom/rroom.html>.
3. *Reactive environments: Throwing away your keyboard and mouse*, 1997. J. R. Cooperstock, S. S. Fels, W. Buxton and K. C. Smith, <http://www.csl.sony.co.jp/person/jer/pub/cacm/cacm.html>.
4. *The Presentation of Self in Everyday Life*, Erving Goffman, Doubleday, Garden City, New York, 1959.
5. "An historical account of the WearComp and WearCam projects developed for personal imaging", Steve Mann, *International Symposium on Wearable Computing*, Cambridge, Massachusetts, October 13-14, 1997, IEEE.
6. "Smart clothing: The wearable computer and wearcam", Steve Mann, *Personal Technologies*, March 1997, Volume 1, Issue 1.
7. "Humanistic intelligence: wearcomp as a new framework for intelligent signal processing", Steve Mann, *Proceedings of the IEEE*, 86(11):2123-2151+cover, November 1998, <http://wearcam.org/ieeecomputer.htm>.
8. "Wearable Wireless Webcam", Steve Mann, 1994, <http://wearcam.org/>.


DATA COMMUNICATIONS PROTOCOLS FOR LINUX

SNA X.25 Bisync HDLC Frame Relay SDLC

Gcom's Board and Software Systems can help you solve your synchronous data communications problems

Available for Intel Pentium® class PCs running LINUX
(Also SCO, UnixWare, Solaris and QNX)

STREAMS-based implementations for maximum flexibility and maintainability from Gcom, Inc., the developer and maintainer of LINUX STREAMS



| | |
|------------------------|---|
| Single Board: | 2-port or 4-port, ISA or PCI |
| Electrical Interfaces: | RS-232, EIA-530, V.35, RS-422, X.21 56K CSU, T1 CSU |
| Speed: | Up to 5Mbps |
| InterSync Controller: | Pentium® PC pre-loaded with LINUX and pre-configured protocol software. Includes interface board. |

Gcom has 20 years of protocol development and support experience. We've helped solve mission-critical problems for VISA, MasterCard, Dow Jones, Unisys and many other successful companies. And we can help you, too.

Visit our web page at
<http://www.gcom.com>

Or call us at
217-351-4241

Quantity pricing available
Reseller inquiries welcome

GC
Om **Gcom, Inc.**
communications software products

1800 Woodfield Drive, Savoy, Illinois 61874
Fax: 217-351-4240 e-mail: sales@gcom.com

Hunting Hurricanes

The authors tell us about hunting hurricanes using the Scanning Radar Altimeter running on Linux and analyzing data with Yorick.

by C. Wayne Wright and Edward J. Walsh

In March 1998, we started development of a new Linux-based data system for the NASA Goddard Space Flight Center scanning radar altimeter (SRA). The goal was to significantly reduce the system weight and volume to enable its installation on one of the NOAA hurricane hunter WP3D aircraft (see Figure 1) for the 1998 hurricane season. The SRA measures hurricane directional wave spectra and storm surge. The data will ultimately be used to help refine and improve hurricane models and improve forecasting and understanding.

The 1998 hurricane season was quite active and the SRA successfully flew in hurricanes Bonnie, Earl and Georges, collecting almost 50 hours of actual mission data.

Our principal obstacle was the short time frame until we needed to be operational onboard the hurricane hunter. The size, weight, complexity and power consumption of the SRA were also critical design items because of floor loading considerations and the limited payload capacity of the P3 aircraft when operating on long (10-hour) missions in turbulent weather conditions (hurricane-eye wall penetrations). Interrupt response time, crash-proofness and freedom from "lock-ups" were all important considerations when choosing the operating system for the SRA.

The new SRA data system, built on top of a Red Hat 4.2 system and Linux kernel 2.0.29, occupies eight inches of vertical rack space, weighs about 40 lbs, runs totally from an internal 12-volt aircraft battery and requires about 120 watts of total input power. It includes a custom ISA board with several PIC microchips which perform dedicated functions

for the radar. It also includes the entire radar IF (intermediate frequency) strip, detectors and a 2ns/point waveform digitizer. No monitor or keyboard is directly connected to the SRA; instead, Linux laptops are used for all control and display. Those laptops run Red Hat 5.1 and 2.0 Linux.

The RT-Linux (Real-time Linux) software does the following:

- Drives the waveform digitizer.
- Computes the centroid-based range measurement between the transmit and return pulses.
- Manages 96 automatic gain control loops.
- Corrects for aircraft attitude and off-nadir angle.
- Deposits formatted data in a shared memory block from which a normal Linux program extracts and records it to a disk file.

The SRA makes extensive use of Tcl/Tk and the Blt graphics library for real-time display.

Post-processing of SRA data is done with Yorick, a free and very powerful programming language that runs on Linux, a wide variety of other UNIX platforms and MS Windows.

Background

The previous implementation of the SRA was developed in 1988 using an array of 68020s on a Multi-bus-I backplane, a CAMAC crate full of nuclear physics instrumentation and a combination of UNIX and VRTX (VRTX is a real-time kernel). VRTX ran on real-time processors and UNIX ran on the system host. The CAMAC crate was quite heavy, consumed considerable power, occupied substantial rack space and was expensive. It used hardware time-interval units (TIUs) to measure the time for a radar pulse to travel from the aircraft to the ocean and back. It used "threshold detection" which caused the TIU to stop and a CAMAC-based waveform digitizer to acquire the return waveform. The waveform data required its own 68020 processor to "process" each waveform and extract certain data. The data were used to refine (post-flight) the range measurement made by the TIU. Threshold TIUs suffer from an effect known as "range walk", which causes the



Figure 1. Front View of NOAA-43, One of Two WP3Ds

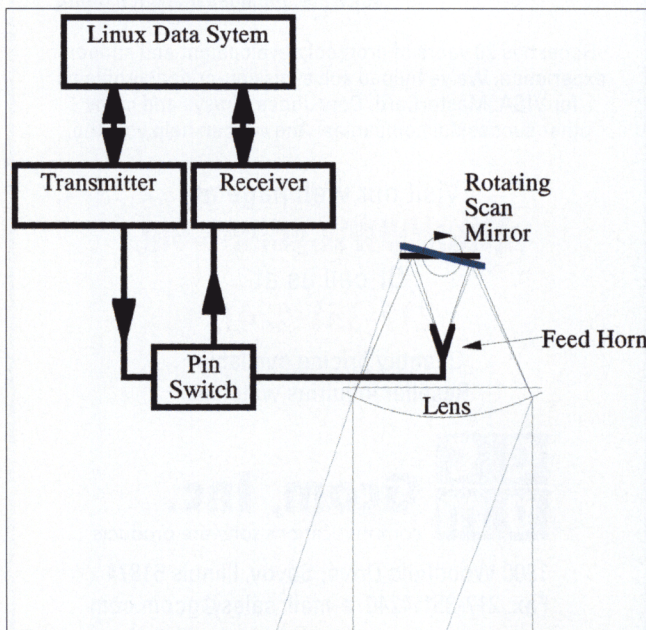



Figure 2. Block Diagram of the SRA Sensor



The most fun you'll ever have being a geek.



thebazaar

THE TECHNOLOGY OF COOPERATION
N E W Y O R K C I T Y



MARCH 13 - 15 1999

JACOB JAVITS CENTER **NEW YORK CITY**

(201) 689-9354 • <http://thebazaar.org> • email: staff@thebazaar.org

The Bazaar is a nonprofit event sponsored by in Limine, an arts and education nonprofit (<http://www.inlimine.org>).

Proceeds will benefit free software projects and development efforts.



measured range to vary as a function of the strength of the return pulse. The array of processors communicated with each other via a 4MB memory card which resided on the multi-bus. Control of the system was via a character-based terminal and real-time display was done on an SBX Matrox graphics module which was managed by its own 68020 processor. One of the 68020 processors ran UNIX; that processor ran programs which extracted radar data from the 4MB card and stored it on a 9-track magnetic tape or a disk file. The UNIX processor hosted all software development and managed the operator control terminal.

Due to its volume, weight and power consumption, we were unable to install this version of the SRA on the hurricane hunter. Limitations in the hardware signal-tracking circuits would frequently falsely trigger the system on a side lobe and effectively eliminate the true range measurement.

SRA System Description

The SRA is an airborne, 36GHz, down-looking, raster-scanning pulsed radar. A simple schematic block diagram of the sensor is shown in Figure 2. Its one-degree beam (two-way) is scanned across the aircraft flight track and a precise time-of-flight measurement is made for each of 64 pulses transmitted at 0.7 degree intervals across the scan. As the aircraft proceeds, a topographic image of the surface (normally ocean waves) is developed, recorded and displayed. The nominal ranging accuracy of the SRA is 10cm. Three differential carrier-phase GPS receivers are used to measure the exact location of three GPS antennas mounted in an array on top of the aircraft. A ground-reference GPS is set up where the flight originates and the ground and aircraft GPS data are processed post-flight to produce an aircraft trajectory, typically accurate to about 30cm in our application. Higher accuracies are possible when operating under less stressful flight conditions.

The Radar Components

The SRA radar consists of a 20-inch Rexalite lens, a feed horn on the lens axis which looks up into a mechanical scanning mirror that mirror-images the feed horn to the focal point of the lens, a pulse modulator and RF exciter, receiver, 1.7KW Extended Interaction Amplifier (EIK) and the RT-Linux data system. The data system is the topic we will discuss here. Figure 3 is a photograph of the SRA scanner installed on the NOAA hurricane hunter. The fairing is removed in this photo.

Figure 4 is a block diagram of the SRA power system. The SRA requires an uninterruptible power source for Linux and the three differential GPS receivers and computers. Instead of an off-the-shelf UPS, we went with a 12-volt 25 AH "RG" (recombinant-gas) sealed aircraft battery as the prime power source for the system. This was chosen for two reasons:

- We needed an uninterruptible power source, because aircraft are notorious for power dropouts during engine start and shutdown.




Figure 3. The SRA Scanner Assembly Mounted on the NOAA P3

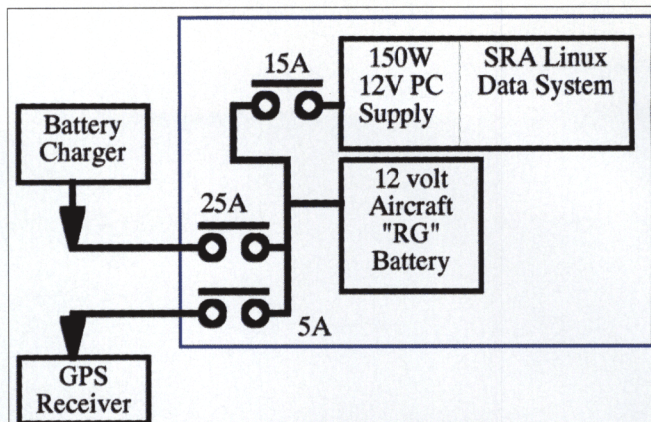


Figure 4. SRA Data Power System

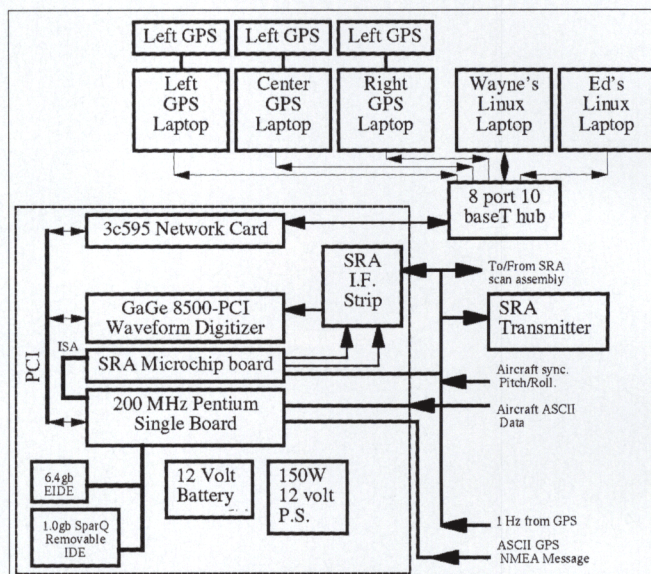


Figure 5. Block Diagram of the SRA Data System

- We needed to power our 12-volt GPS receivers for up to an hour before and after each mission without aircraft power applied.

We purchased a 12-volt input 150W PC power supply to power the data system. The battery can power the data system and the three GPS receivers for about two hours, or the GPS receivers alone for five hours. We located the battery in the rear of the custom data system housing. Figure 4 depicts the wiring of our power system.

Figure 5 is a block diagram depicting the internals of the SRA data system. The computer is a single-board 200 MHz Pentium which plugs into a passive backplane with ISA and PCI slots. The CPU card contains PCI-VGA video, PCI-IDE controller, PCI fast-wide SCSI controller, 64MB of RAM, 512MB of cache, two serial ports, a parallel port and the CPU. A PCI 3c595 network card provides networking and a special-purpose ISA card loaded with PIC microcontrollers provides an interface to the radar systems. A 6.4GB EIDE disk drive is used as /dev/hda to hold Linux and for data storage. A backup SparQ 1.0GB removable drive is installed as /dev/hdb. The system has no floppy or CD-ROM drive. If a CD or floppy is needed, they are simply remotely mounted with NFS from one of the Linux laptops which have both. No keyboard or monitor is used for normal operations, though they can be plugged in if the need arises.

Initially, we used a 4.2GB SCSI drive, but that used too much electrical power. Early development was done using a 250W 117vac PC power supply. When we switched to the 12-volt input 150W power supply, we discovered we were over our power budget by 25 watts or so. During the boot process, the power consumed by the combination of the SCSI drive and the waveform digitizer would cause the power supply to "spike" the 5-volt source and cause a reboot. It took us several hours to find this problem. It would generally happen just as Linux began loading, due to the digitizer being powered up and the drive being accessed. During the DOS boot, the digitizer was not powered until after DOS booted and after the digitizer configuration program loaded and ran. Consequently, the loading of Linux was the straw that broke the camel's back. We finally settled on a 6.4GB EIDE disk drive for Linux and for data storage. The power consumption of the EIDE drive is substantially less than the SCSI and no perceptible difference is seen in performance of the data system.

Figure 6 is a photo of the SRA data system during development. It was in this "state" until just a few days before our first test flight on the NASA C-130. Figure 8 is a photo of the data system after packaging. Figure 7 shows the internal organization of the data system as viewed from the top rear. The enclosure is reversed from most rack mounts. We wanted to have ready access to the computer card connections without having to remove the rear rack cover. The only connections on the rear are for the GPS receivers and the battery charger. The black power supply under the data system in Figure 8 is our prime power supply/battery charger.

2ns PCI Waveform Digitizer, DOS, GageScope

The core data acquisition device in the SRA is the GageScope 8500-PCI waveform digitizer. It provides for up to 128KB of sequential samples taken every 2ns (nanoseconds). This permits us to digitize a 256-microsecond waveform. We actually digitize for 60 μ s, beginning a few hundred nanoseconds before the radar pulse is transmitted and ending after enough time has expired to accommodate a signal return from our highest possible altitude. The pulse takes 2 μ s to travel 1000 feet and return, so to accommodate a maximum altitude of 30,000 feet, we need to digitize at least 60 μ s. Since a point is digitized every 2ns, there will be 30,000 points in each waveform. We don't read all 30,000 points out of the digitizer. We "track" the position of the returns and




Figure 6. The SRA Data System during Development




Figure 7. View of SRA Data System Internals





Figure 8. Data System after Packaging



Embedded Empress RDBMS delivers power, speed and efficiency to embedded applications.

Fast and Compact. The RDBMS engine requires about 1MB of RAM. Optimized for performance, the Empress RDBMS is just the right size for your applications.

Easy to Use. Interactive dialogues and simple UNIX/WIN and ANSI SQL standard commands make the embedding of your Embedded Empress RDBMS a snap.

Priced Right. The Embedded Empress Developer's Toolkit costs a fraction of the price of comparable systems from Oracle, Sybase, or Informix.



FREE EVALUATION

US: 301-220-1919 Int'l: 905-513-8888

www.empress.com
info.linux@empress.com

read out only 256 points centered around where we expect the return to come from. Since the ocean is basically flat, this technique works well.

The driver code provided by Gage for the 8500 supports DOS, Windows and Windows NT. It is extensive, to say the least. It contains several thousand lines of code solely to initialize most of the cards that Gage makes to an operational state. Apparently, much of the functionality of the card is loaded into programmable Logic Arrays from the DOS driver. The Gage driver code supports virtually every waveform digitizer made on several different OS platforms. They make extensive use of conditional compilation to select both the desired digitizer board and the desired operating system. They attempt to establish an isolating layer of driver code, so that a common set of driver calls appears to the users of their supplied library.

After looking at the driver start-up code, we thought it might take more time to port the start-up code to Linux than we could afford. In order to avoid porting the long and complex start-up code, we elected to make the system dual boot Linux and DOS. This scenario has worked well, permitting us to get a DOS program going quickly which would configure the digitizer. After the digitizer is configured, the autoexec.bat DOS script loads Linux using **loadlin**, a DOS program which can load a Linux kernel from DOS. The DOS digitizer start-up code leaves the digitizer in a known functional state. The code required to use the digitizer is actually not very extensive and only requires accessing a few registers and memory locations on the Gage card. The folks at Gage were very helpful in getting it working.

Hardware Interrupts

Waveform data are extracted from the digitizer after a radar pulse event has occurred. One of the 16C65A microcontrollers controls all aspects of triggering the transmitter, actuating various gates, triggering the waveform digitizer and finally interrupting the Linux waveform digitizer interrupt handler.

The RT-Linux interrupt typically responds in $2\mu\text{s}$ (on our 200MHz Pentium) with occasional jitter to several microseconds. When we did this same test with MS Windows a couple of years ago, we found the fastest response to be on the order of $50\mu\text{s}$ (486-dx2 66MHz) with jitter well into tens of milliseconds. It is incredible just how responsive Linux is to interrupts. Figure 9 is a digital scope capture, where the top trace rising edge is the actual hardware interrupt signal on the ISA backplane. The bottom trace is a hardware signal generated by the interrupt code. It simply wrote a "1", waited awhile, then wrote a "0" to the printer port. Each horizontal division is $2\mu\text{s}$. This demonstrates the typical latency of our RT-Linux system. Concurrent with this test, we ran a **find** command in another xterm so the system had something to do.




Figure 9. RT-Linux Interrupt Jitter

Microcontroller Board, 16C65A

Microcontrollers permit very hard and reliable real-time capabilities. They are well-suited to replacing arrays of chips and digital logic in many applications such as the SRA. We designed a special ISA interface board for the SRA which encompasses most of the special requirements of the radar and special interfaces.

The board is presently populated with four Microchip 16C65A microcontrollers. One microcontroller implements a real-time clock which automatically maintains time-of-day synchronization with our GPS receivers. It has a least significant fractional time bit of 200 nanoseconds and provides the SRA with up to 64 bits of accurate time information. This chip automatically captures the trigger time for each radar pulse. It and its neighbors can all be read and written by Linux.

A pair of microchips function together to convert the scan encoder pulse trains into radar trigger events. As our scan mirror rotates, a scan encoder measures the scan angle. At our scan rates, it produces a pair of 40KHz square-wave signals which are 90 degrees out of phase. One microchip is programmed to combine these two signals together and produce a single 80KHz signal, which is then counted to determine the position of the scanner. The second microchip is programmed to count the 80KHz signal and initiate a radar pulse at predefined angles. With its 200ns instruction time, this microchip directly controls all aspects of the transmitter and receiver electronics and also generates an interrupt to Linux once a waveform has been acquired by the waveform digitizer. For each SRA pulse, this microchip:

- Protects the receiver front-end from damage.
- Verifies that the receiver is protected.
- Gates the digitizer on.
- Gates the EIK amplifier on.
- Delays exactly 200ns.
- Triggers the transmitter modulator to generate an 8ns pulse and causes the real-time clock microchip to capture the present time.
- Delays 200ns for the transmit pulse to be well clear.
- Enables the receiver to receive return signals.
- Interrupts the RT-Linux SRA module to extract the waveform data from the digitizer.

RT-Linux

RT-Linux is a patch which gives Linux many of the most important features needed by real-time programmers and embedded-system designers. It is implemented as a set of modules which can be installed and removed using **insmod** and company. You also use **insmod** to install any real-time code you write. RT-Linux programs execute in the kernel space and have full access to the system hardware and kernel code as well.

We've done a considerable amount of development using Turbo-C and DOS in the past, and it is truly amazing how infrequently we had to reboot Linux during development of the SRA. Back under DOS, we usually had to reboot several times per day. With Linux, we had to reboot only three or four times during the entire development period.

Shared Memory

Once the RT-Linux programs/modules capture the data, they must be written to storage and displayed for the system operator. We accomplish this by using shared memory. The SRA has 64MB of RAM and we configured the kernel to boot using `mem=61m` which causes the kernel to manage only the lower 61MB, leaving 3MB untouched. It is this 3MB that we use for real-time data capture and as a common communication buffer area between RT-Linux modules and normal user-space programs. Figure 10 depicts the SRA memory usage.




Figure 10. SRA Memory Usage

We wrote a single C program (`rgc.c`) which provides most of the interface between Linux user mode and RT-Linux. This program is a simple command-line style program with tons of commands to read and write data space in common between RT-Linux and user space. Most of our Tcl/Tk scripts merely open a pipe to this program and use it to pass commands and extract data from the system. The

program can also be used directly from the command line. This makes development and debugging simpler.

One of the run-line options to `rgc` causes it to loop, testing for data to be written to disk. If no data are ready, the program sleeps for one second. If data are ready, they are extracted and written to the specified disk file.

Linux Laptops

We use up to five laptops on the SRA at once: three for collecting GPS data (one laptop for each GPS receiver) and two for control and display of real-time SRA data. A personal laptop is used for control, and if we're both on the flight, we can both run several instances of the same display programs using another personal laptop. We each have our favorite color-bar for the image of the sea. We'll frequently use one machine to control the SRA and the other to write or modify display or system software as we're flying. The laptops are Chembook 9780s. Each has a 4GB internal hard drive and a modular 6.4GB drive (in place of the floppy), a 14.2" XGA LCD display, PCMCIA Ethernet card and a 233MHz Pentium-Pro CPU.

Each of these machines dual boots either Red Hat Linux 5.1 or MS Windows 95. To use the laptops as X terminals, we boot Linux, then run the Xfree86 server. We run the X server such that the laptop becomes an X terminal for the SRA data system. This puts most of the burdensome display processing on the laptop processor, since the X server seems to be where the CPU cycles go. There are two ways to cause X to act as an X terminal. The first is:

```
X -query somemachine
```

Backup your network from only \$199

Special Introductory Price for Linux

arkeia⁴⁰



- ▶ High performance Network backup ... 200 clients at a time; 6 to 60 GB/hour.
- ▶ Servers: Linux, Solaris, SGI, AIX, Alpha (osf 1), HP/UX and SCO
- ▶ Clients: Linux, Windows 95, Windows NT, Solaris, SGI, AIX, Alpha (osf 1), HP/UX, SCO, OS/2, Novell, VMS, ICL, and more
- ▶ Drives: 4mm DAT including DDS3, 8mm Exabyte, DLT, Magstar, DTF Sony.
- ▶ Libraries : ADIC, ATL, Exabyte, HP, IBM, Overland, Storagetek, Tandberg, and more.
- ▶ Automated backup scheduling
- ▶ Shareware version available for home use



For more information contact: sales@knox-software.com

Knox Software USA
(888) 333 - 5669

Knox Software Europe
+33 1 53 09 99 99

Plus Tax and Shipping & handling

30-day free demo
www.knox-software.com

and the second:

```
X -indirect somemachine
```

The target machine must be running XDM (X display manager) for this to work. The first method will link directly to the target machine, where you see a typical XDM login prompt. This first method is what we use when controlling the SRA data system. The second method will give you a list of all the machines known on the network to support XDM or X terminals. It is useful back at the lab where many potential hosts are available to pick from.

You can even have two or more X servers running at once. Here's an example:

```
X -query first-machine
X :1 -query second-machine
X :2 -query third-machine
X :3 -query fourth-machine
```

You can get a local X server going with the command:

```
startx -- :4
```

The SRA system configured for storm-surge measurements consists of three Chembook Pentium laptops which dual boot Linux and DOS. The GPS data acquisition program was written for DOS, so each laptop runs this DOS program when collecting the GPS data. After the mission, we reboot the machines to Linux and transfer the data to the SRA data system where it is archived with the other mission data. Once it is all together, we transfer it to the two laptops. In this way, we have triplicated the data. We then take the laptops with us back to the hotel and begin analyzing the data. All five laptops and the SRA data system are on a 10baseT Ethernet network.

GPS and RS-232 Aircraft Data

Some aircraft data are read via RS-232. For this, we are using the standard `/dev/ttySxx` ports and drivers. The aircraft data are in a 9600 baud stream occurring once per second and the GPS produces a position message twice per second. We use our GPS message to drive a simple Blt plot of the latitude versus longitude, so we can track the progress of the flight.

The RS-232 data are actually captured by the `rgc` program, since the RT-Linux modules can't make use of the native Linux drivers and we didn't need to rewrite drivers that were working perfectly. Once the data are read, they are copied to the shared memory area above 61MB where any of the programs can access it. Normally, it is accessed by another invocation of `rgc` and `read`.

Pitch, Roll, Heading and Track Angle

Accurate aircraft attitude, heading and track angle data are critically important to the SRA in real time. The pitch-and-roll attitude of the aircraft is taken from the on-board Inertial Navigation Units (INU), using Synchro-to-digital converters—one for each parameter. These are read by the RT-Linux module during each scan line. The heading and track information is presently provided via RS-232 from the on-board aircraft data system, which has a direct interface to the INU's digital data stream.

Resulting Data (Radar, GPS, Aircraft)

The SRA radar data are written to disk files by the `rgc` program. The aircraft data are captured by a separate pro-

gram and written to a separate disk file. This data is normally captured for the entire duration of the flight, providing a complete flight record in a single file. The carrier-phase GPS data are captured continuously from 45 minutes before the flight until 45 minutes after the flight. The pre- and post-mission data are necessary to resolve the aircraft position to the centimeter level.

System Software Development

Before any Linux development was carried out, we felt it necessary to write some DOS code to work with the Gage digitizer board. Turbo-C version 5.0 was required to compile and use the Gage-supplied library. Once we were successful in getting a Gage example program to work on DOS, we worked with Gage Engineers to communicate directly with the digitizer using a normal user-mode program. The main trick was to make the DOS program configure the digitizer and then exit without powering it down. The second trick was to boot from DOS into Linux; this turned out to be quite easy with `loadlin`.

We determined the PCI board settings for the digitizer by reading `/proc/pci` and then hard-coding various test programs with the values. We wrote various normal user-mode programs to become familiar with the digitizer. We were able to manipulate the digitizer card in every way except handling interrupts. The `gdb` debugger was a big help throughout the development.

Microcontroller Software Development

A substantial part of the SRA software is actually firmware resident on various microchips.

Microchip provides, at no cost, a very complete and easy-to-use development package for their 16C65A (and other) microcontrollers. It sports a comprehensive simulator, making it possible to watch simulated execution of quite extensive programs. The only downside is the system runs only on MS Windows.


RT-Linux

The RT-Linux extensions provide just the right features for a real-time data system such as the SRA. The extensions provide much more capability than we actually use in the SRA. We use it to start an RT task at the end of each raster scan. The task processes all the data captured during the previous scan and makes a number of calculations necessary to configure the system for the next raster.

Linux User to RT-Linux Interface

We wrote `rgc.c` to be a liaison between normal user processes under Linux and the RT-Linux SRA module. Quite simply, `rgc` sets up a pointer to the shared memory space that the SRA RT module uses for data storage. They understand each other because they share a common `.h` file defining the data organization in the shared memory space. Figure 11 depicts how the various SRA real-time programs communicate with each other. `rgc` usually reads commands from `stdin` and writes to `stdout`. If it is invoked with certain switches, it forks and polls for RS-232 data and/or writes captured data from the shared memory to disk, all the while taking its commands from `stdin`. The command set is simple ASCII strings such as `set thresh 24` or `get roll`. The `Tcl/Tk` programs each open a pipe to their own private `rgc`, then send commands and receive data back. Everything is done this way except the topographical image display. That program,

A DIFFERENT KIND of Animal



A Different Kind of Solution

O'REILLY
www.oreilly.com

101 MORRIS STREET, SEBASTOPOL, CA 95472

ORDERS/INQUIRIES: **800-998-9938**

WEEKDAYS 6AM-5PM PST

707-829-0515 • FAX: 707-829-0104

EMAIL FOR OUR CATALOG:

CATALOG@OREILLY.COM

INCLUDE YOUR NAME AND MAILING ADDRESS







O'REILLY BOOKS ARE ALSO AVAILABLE
AT YOUR BOOKSTORE

PLEASE MENTION CODE **ALXJ** WHEN ORDERING

O'Reilly creates books for professionals like you — people with serious information needs who don't want run-of-the-mill answers.

Our readers depend on us to deliver no-nonsense, reliable solutions to their technical problems. You'll never find any ten-pound doorstops that have been rushed into print here. Instead, you can count on us to produce meticulously researched, authoritative books written by the experts — with an unbiased, independent point of view.

Our readers know that. It's why they trust us.



creep.c (because it creeps up the screen), accesses the shared memory directly. The main reason for concentrating everything into rgc is that it generally means we need only recompile rgc, creep and the SRA module when something is added or removed in the shared memory area. In short, it makes for quicker development.

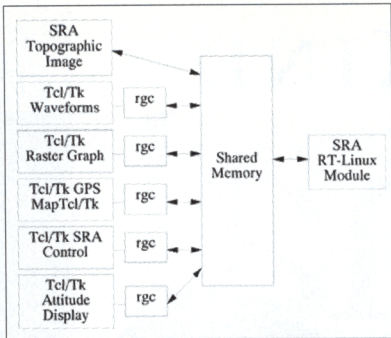


Figure 11. SRA Program Block Diagram

Linux X Terminal—System Display and Control

Figure 12 is a screen shot of the SRA control laptop during hurricane Bonnie's landfall. The image on the left side of the screen is the real-time topographic display. It is grayscale encoded so that the higher things are, the more white they appear; the lower, the darker. This image clearly shows waves on the left side of the image, the beach in the center and a very distinct dune line. We also have a color-encoded version of this program, but its interpretation is not as intuitive. The blue/brown display represents the attitude of the aircraft. It is a short Tcl/Tk script which reads aircraft attitude data captured by the SRA RT-Linux module.

The bright green display shows how we control and designate the operating conditions for the SRA. At this

time, we manually find the return signal using the slider. Once found, we click the "auto" button and the system will keep the ground in the center of our digitizer window, regardless of aircraft altitude variations. The flight map is yet another short Tcl/Tk program. It extracts GPS position data from the shared memory area and uses it to map our position.

TK, Blt, Xview

Tcl 7.6 and Tk 4.2 with Blt 2.3 are used extensively in the SRA. Initially, we thought it might be useful only for prototyping, but it soon became obvious that the X server would be the display bottleneck and not Tcl/Tk.

During development and before we purchased the laptops for control, we used a monitor connected directly to the SRA system. This meant that the X server would run there too. When we began experimenting with using a remote X server, we quickly discovered that the burden of the X server had also moved to the remote system. This was a no-effort way to automatically distribute the load across one or more computers in the system.

We wrote the image display in C using the Xview library. We used this library because we already had a book about it, and it didn't look too difficult to use. It writes each scan line directly to the display and simultaneously to a "pix-map". When a "repaint" event occurs, the pix-map is used to repaint the whole image. A great way to put a load on the display computer X server is to grab the image map and move it around the screen. The load on the displaying computer will go through the roof, but the data system will remain unaffected.

Data Analysis—Yorick

Once we had some SRA data, we obviously needed to build some software to review it. We wanted to have SRA processing software on several machines and without licensing hassles. That way, we would be able to develop programs at home, on an office laptop (which is also used to control the SRA), on the SRA data system computer and on office Linux and Windows PCs. In total, we needed processing on at least five to ten different systems. We considered IDL, Matlab and Yorick. Our tool of choice for processing was Yorick. It is free, very powerful and will run on a wide variety of platforms including almost every UNIX machine known, Linux and Windows. It has the ability to save data so

DESIGN HAS NEVER BEEN EASIER!

ENVIRONMENT: LINUX, WIN 95/98/NT

COMPLETE CREATION OF 2D DOCUMENTATION

3D-SOLID MODELING

CALCULATIONS OF MECHANICAL COMPONENTS

AUTOMATIC CREATION OF BILL OF MATERIAL

OPTION TO CONNECT TO AN INFORMATION SYSTEM

THE BEST PERFORMANCE-TO-PRICE RATIO

E-mail: mail@varicad.com
<http://www.varicad.com>

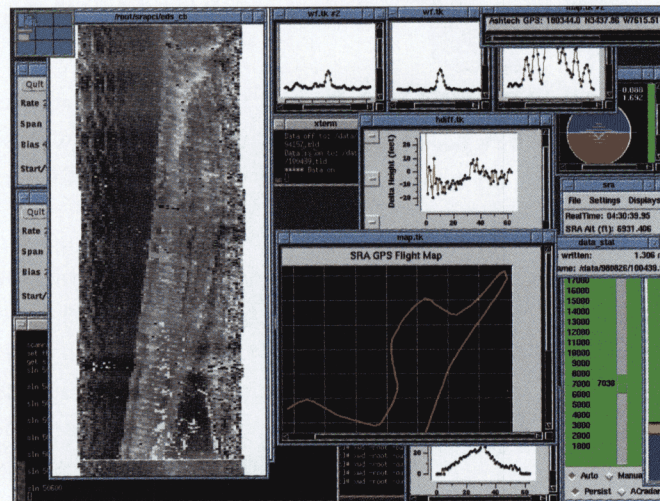


Figure 12. SRA Screen in Operation during Bonnie.

it can be read on a big-endian or little-endian machine.

We first heard of Yorick from an article in *Linux Journal* ("The Yorick Programming Language", Cary O'Brien, July 1998). We downloaded it and gave it a try. We like its C-like syntax and ability to load (and reload) individual functions of a program. This makes for a very powerful and flexible development environment. One of its best features is its cost—free! To put either Matlab or IDL on all the machines would have been prohibitively expensive. Since we have Yorick on the SRA data system and on the controlling laptop, we can easily analyze data in the field with minimal effort using the Linux laptop. Figure 13 shows topographic images from the SRA overlaid on a wind field plot from the August 24th flight. The sea state was above 18 meters (60 feet) on the northern flight line.

Results

We had two or three short test flights on the NASA C-130 aircraft before we had to pack everything up and ship it to MacDill Air Force Base in Tampa, Florida, for installation on the NOAA hurricane hunter. We removed a number of bugs during these test flights, but not all. When we shipped the system, it still would not track properly.

Once we were all installed on the hurricane hunter, we had a 6-hour test flight. This permitted us to work out almost all of the bugs we had seen earlier and a few new ones. We still had a few problems with the tracking code, which would not track reliably.




Figure 13. Initial Data Product from Yorick Showing Surface Topographic Images Superimposed on NOAA Wind Plots

Fast Linux ISAM Access

Linux Support Since 1995!

**Puzzled by
database
performance
issues?**

**Nothing
beats... ISAM
performance**

Real-world data management solutions are typically more complex when one examines the pieces, than initially recognized by the majority of database programmers. All software projects are complex puzzles comprised of many details, most of which are data-related. Often today's "DBMS" solutions sacrifice the speed or control essential for a competitive application.

c-tree Plus®, by FairCom, has been the choice of commercial developers for twenty years precisely because it offers the **flexibility** and **control** at the detail level to fit a wide variety of data management needs.



Proven on large Unix servers and workstations, c-tree Plus's small footprint and exceptional performance have also made it the engine of choice for professional developers on Linux, Windows and Mac. c-tree Plus offers sophisticated ISAM level control with which the developer may define precise data management solutions, making it a perfect fit for any development project requiring specific data handling features.

**c-tree Plus® offers the most
mature ISAM solution today...**

FairCom's c-tree Plus database engine:

- Advanced Indexing Technology
- **Complete Source Code, just like Linux**
- Complete Transaction Processing
- ODBC Interface
- Over 25 Developer's Servers Included
- Portable Multi-threaded API
- **Royalty Free**
- Standalone, Multi-user or Client/Server Models
- Thread-Safe Libraries
- Y2K Compliant

The FairCom Server:

A solid, high performance database server that is scalable, portable and offers unequalled control. FairCom has been providing database solutions to the commercial development community for twenty years. You won't find a better Linux DBMS solution, with these features and performance anywhere else!

- Client Side Source Code, **just like Linux**
- File Encryption
- File Mirroring Logic
- Full Conditional Index Support
- Full Heterogeneous Networking
- Multiple Communication Protocols
- Online Data Backup
- Small Memory Footprint
- Flexible OEM Licensing Options
- Source Code Availability



All these platforms supported in one package:

Linux (Alpha...), MIPS ABI, DEC Alpha, Sun SPARC, Windows 9X, SCO, 88OPEN, AIX, RS/6000, HP9000, Sun OS, Interactive Unix, AT&T System V, QNX, Free BSD, OS/2, Mac, Windows NT, Windows 3.1, DOS, Netware NLM, & Banyan VINES.

Support for most major Linux versions, including:

- Caldera
- Debian
- Red Hat
- Slackware
- S.u.S.E.

Don't wait, see for yourself!
USA. 800.234.8180

www.faircom.com
Email: info@faircom.com

FairCom®

Database SOLUTIONS Since 1979

Phone: USA 573.445.6833 - EUROPE +39.035.773.464 - JAPAN +81.59.229.7504 - BRAZIL +55.11.3872.9802
Other company, product and operating platform names are registered trademarks or trademarks of their respective owners.

Bonnie

We flew two missions in hurricane Bonnie: the first on August 24 and the second during landfall on August 26, 1998. During our first transit flight from Tampa to the storm, we were able to isolate and correct the tracker bug and everything started working better than expected. Soon after leaving the east coast of Florida, our topographic display of the sea came alive for the first time, showing real sea state. Ocean waves as high as 63 feet were observed in the northeast quadrant of the hurricane on the 24th. Figure 14 shows our August 26 flight track during landfall overlaid on the aircraft weather radar image and a contour plot of the wind field data. The base image includes the weather radar, the wind field and the coastline and was provided by the Hurricane Research Division (HRD) of the NOAA Atlantic Oceanographic and Meteorological Laboratory (AOML) in Miami. We produced this overlay using Yorick.

In addition to hurricane Bonnie, we also flew in Earl and Georges.

Conclusion

Thanks to the reliability of Linux and all of the off-the-shelf real-time data processing programs available in that domain, we were able to put together a state-of-the-art data system on a very tight schedule with a great variety of real-time displays. The displays proved to be of great value both in troubleshooting during development and in real-time




Figure 14. Flight Track during Bonnie's Landfall

geophysical assessment and interpretation during data acquisition. As a result, we were able to document for the first time the spatial variation of the wave field in the vicinity of a hurricane and the spatial and temporal variation of the storm surge associated with hurricanes on landfall.



C. Wayne Wright (wright@osb.wff.nasa.gov) is a Data Systems Engineer for the NASA Goddard Space Flight Center, Laboratory for Hydrospheric Processes, Observational Sciences Branch, Wallops Island, VA. He is a 1984 graduate

of the University of Maryland with a degree in Computer Science. His interests include aviation, amateur radio and computers. Away from work, he and his wife Vicki operate a Linux web server.



Edward J. Walsh (walsh@osb.wff.nasa.gov) is a scientist for the NASA Goddard Space Flight Center, Laboratory for Hydrospheric Processes, Observational Sciences Branch, Wallops Island, VA. He received B.S. and Ph.D.

degrees in Electrical Engineering from Northeastern University in 1963 and 1967, respectively. Ed is presently on assignment for NASA at the NOAA Environmental Technology Laboratory in Boulder, Colorado.

Resources

"Airborne remote sensing applications to coastal wave research", Paul A. Hwang, Edward J. Walsh, William B. Krabill, Robert N. Swift, Serdar S. Manizade, John F. Scott, Marshall D. Earle, *Journal of Geophysical Research*, Vol 103, No. C9, Pages 18,791-18,800, August 15, 1998.

"Directional wave spectra measured with the surface contour radar", E. J. Walsh, D. W. Hancock, D. E. Hines, R. N. Swift, and J. F. Scott, *Journal of Physical Oceanography*, 15, 566-592, 1985.

"An observation of the directional wave spectrum evolution from shoreline to fully developed", E. J. Walsh, D. W. Hancock, D. E. Hines, R. N. Swift, and J. F. Scott, *Journal of Physical Oceanography*, 19, 670-690, 1989.

NASA: <http://AOL11.wff.nasa.gov/>

osb: <http://osb.wff.nasa.gov/>

SRA home page: <http://AOL11.wff.nasa.gov/sra/>

hrd: <http://www.aoml.noaa.gov/hrd/>

aoc: <http://www.nc.noaa.gov/aoc.html>

Gage: <http://www.gage-applied.com/>

microchip: <http://www.microchip.com/>

yorick: <ftp://ftp-icf.llnl.gov/pub/Yorick/yorick-ad.html>

Tcl/Tk: <http://www.tcltk.com/>

Blt: <http://www.tcltk.com/blt/>

rt-linux: <http://luz.cs.nmt.edu/~rtl原因/>

chembook: <http://www.chembook.chemusa.com/>

ashtech: <http://www.ashtech.com/>

Concorde Battery: <http://www.concordebattery.com/>

Red Hat: <http://www.redhat.com/>

CTAR
BUSINESS CRITICAL BACKUP

- Fastest Backups Available For LINUX
- Quick File Restores
- X-Window & Character Interface Available
- Network Enabled
- Full Auto-Changer/Tape Library Support

For More Information: www.unitrends.com

Or Call: **1-800-648-CTAR**

Outside US: (843) 626-2878
Email: sales@unitrends.com
FAX: (843) 626-5202

Unitrends
"The Developers of CTAR"

Developing your CODE

[Shouldn't feel like a free fall.]

Lock into the stability of Cygnus
GNUPro® Toolkit for Linux.



Key Features

- ANSI-C conformant compiler
- ANSI compliant C++ compiler
- GNUPro visual debugger
- Macro-assembler
- Linker/loader
- Make/build utility
- C and C++ libraries
- Binary utilities
- Printed and online documentation
- Installation CD
- For Red Hat® Linux 5.1 and 4.2

You've been there before. Trying to compile complex code with less than satisfactory tools. You need a professional compiler and debugger that you can trust. Tools that have been tested, proven reliable, and specifically designed for your Linux development environment. You need Cygnus GNUPro Toolkit for Linux.

Take advantage of the GNUPro Toolkit's:

• Robust and tested tools.

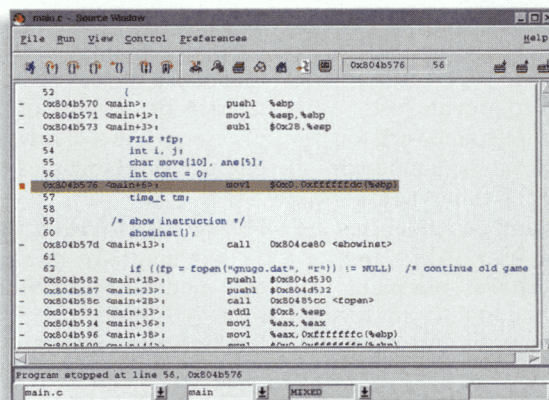
The ANSI-conforming C and C++ compilers, macro-assembler, visual debugger, utilities, and C/C++ libraries have been rigorously tested and enhanced to support development on Linux.

• Intuitive visual debugger.

The GNUPro visual debugger helps you to quickly visualize, locate, and debug code at the source and assembly level.

• **Easy installation and complete documentation.** With GNUPro Toolkit, you receive an easy-to-install CD, online and printed guides, and a quick reference card to help you become immediately productive. And, to give you a running start, you receive 30 days of installation support.

• **Open source.** Because GNUPro is based on open source technology, you're never locked into a dead-end proprietary software product.



GET IT \$79

Certified

OFFICIAL
CYGNUS
GNUPRO
TOOLKIT

Cygnus has
more GNU
experts than
any other
company in
the world.

Make sure that your GNU
tools are Certified Official
from Cygnus.

Need extra assistance?
Upgrade to Cygnus
Mission Critical Support.
A variety of programs
and options are available.

Visit our Web site at
www.cygnus.com/gnupro/lj

Call us at
+1 800 CYGNUS 1



Copyright © 1997-1998 Cygnus Solutions. Cygnus GNUPro Toolkit is a trademark of Cygnus Solutions. All other trademarks are the property of their respective owners.

COAS: A Flexible Approach to System Administration Tools

Caldera is working on a new easy-to-use configuration tool for Linux. Mr. Kirch gives us the details.

by Olaf Kirch

COAS stands for Caldera Open Administration System. It will be incorporated as the main configuration tool in future versions of the OpenLinux distribution.

For those who have never used OpenLinux, the tool we have been using for quite a while is called LISA (Linux Installation and System Administration), which is basically one huge shell script using a modified version of the **dialog** tool to interact with the user. When we felt it was time to move on to something new, we of course looked at what was already available. The only viable option at that time seemed to be LinuxConf, which had quite a ways to go before it would become useful. Since that time it has become much better, but because we had already started work on COAS, we decided to stick with it. Of course, we believe our concept is better.

The source code to COAS is released under the GNU General Public License. We feel our work might be useful to the Linux community as a whole and we want to invite interested programmers, administrators and users to participate in its development by offering comments contributing patches or even modules.

Vertical Modularity

The main idea behind COAS is not to provide just another administration tool, but an entire framework for writing one. From the start, we wanted it to be a modular application where assumptions about such things as system data representation, file locations and dependencies are separated as much as possible from user dialogs and vice versa. Ambitious as this goal may appear, our main interest was the ability to easily adapt the tool to changes in the underlying platform and in porting it to other Linux platforms.

I like to call this vertical modularity, because it breaks up the task of system administration into three layers. At the lowest level are native system data files, such as `/etc/passwd`, `/etc/hosts` or files that define the IP address for a particular network interface.

On top of that, COAS implements an internal representation as a kind of database. If this term made you jump in your seat and shout, "Oh no, Mr. Bill, not a Linux Registry!", please be assured that this is definitely not what we want it to be. COAS is supposed to be vi-administrator friendly. We want users to be able to switch between COAS and vi (or Emacs) administration, because even though we hope COAS will be useful for everyday tasks, it cannot cover each and every feature of a system component. (Consider the configuration monster incarnate, **sendmail**—you can spend as much time writing configuration software for it as Eric Allman keeps churning out new features.)

The native system files will remain the primary source of information. The COAS data model is strictly a run-time representation of system data that attempts to hide the on-disk representation from the upper layers. For instance, an administration module for the BIND server should not have to bother about where DNS zone files are located and how they are to be parsed; all it needs is the list of DNS zones this server is a primary or secondary name server for and the records they contain.

Having an abstract data representation also allows for alternate data access mechanisms. For example, our database engine can store a change log of an administration session to a file, which could then be distributed to other machines, thus allowing for bulk updates. Also, there's the vague idea that COAS might one day support remote access via LDAP or SNMP.

The top-most layer is the user interaction code. This code drives the dialog with the user and controls what information is displayed to the user at what time. It uses a standard set of dialogs, provides on-line help, etc. We decided to use a scripting language, Python, at this layer in order to allow for rapid prototyping. In addition to this, wrapping all lower-level functionality in Python classes and functions provides an additional level of insulation that restricts the number of tricks a programmer can pull. This may seem like a disadvantage to the hackers among you, but it is truly a big plus when it comes to code maintenance.

Horizontal Modularity

You may have guessed from my choice of the term "vertical modularity" that there is also a horizontal one, and so there is. Consider the following scenario: a security problem or other misfeature requires you to update a component of your system, such as the BIND name server. Alas, the update is from version 4.9 to version 8.2, which uses an entirely different configuration file format. We could now ask you to install an all-new version of our administration tool in order to accommodate the new configuration file format. On one hand, that is costly in terms of bandwidth. On the other hand, making sure the tool operates properly with all possible combinations of updates applied or not applied would be rather time-consuming for us. The ideal solution would be to package the DNS server administration module alongside our BIND update.

We are attempting to accomplish the following: COAS lets you rip out an entire module, including the data model definition, Python code, message catalogs and so on, and replace it with a different version. We have nicknamed these CLAMs, which is short for Caldera Linux Administration Module (we invented the acronym first and then decided on its meaning, in case you were wondering).

OpenLinux™—a full solution you can sink your teeth into.



LINUX™

(kernel and packaging)



OPEN LINUX™

LINUX FOR BUSINESS

Integrated office suite from StarDivision
with word processor, presentation software,
spreadsheet, E-mail, and Internet access

GUI

Easy install

Includes applications from:

▪ Novell ▪ Sun ▪ Sybase ▪ Netscape

Tested, Proven, Supported

Training and Technical Support Programs

A COMPLETE BUSINESS SOLUTION



To buy OpenLinux™ through a reseller near you, call 1-888-GO-LINUX or visit www.caldera.com.



Data Model

Let's take a closer look at the internal data representation. All information is stored in a tree, with each node having an individual name. For instance, the node containing the password of the root user is named `system.accounts.users.0.password`. If you're familiar with SNMP, think of the way SNMP variables are named.

Nodes can have different types; e.g., `system` is a directory, `users` is a list and `password` is a scalar. Scalar nodes can have various constraints attached to them; for instance, a string may be required to match a regular expression or contain only values from a predefined set of choices. You can also attach your own parsing and representation functions (written in C) to a scalar type, creating custom types that do such things as convert date strings, e.g., *Jun 12* or *tomorrow*, to UNIX time.

All this information is provided by the so-called schema. The schema acts as a sort of blueprint for the data model in much the same way an SNMP MIB definition describes the types and organization of entities for SNMP.

For instance, the definition of the mouse parameters might look like this:

```
MODULE      "PERIPHERALS/MOUSE"
MSGCATALOG "peripherals/mouse"
TYPEDEF DevicePathName STRING MATCHES
    "/dev/[a-z0-9]*"
TYPEDEF MouseProtocol STRING IN CHOICE {
    "Busmouse", "MouseSystems",
    "Mouseman", "Microsoft"
    ...
}
device      RECORD {
    model     STRING
    protocol   MouseProtocol
    deviceFile DevicePathName DEFAULT
        "/dev/mouse"
    emulate3btn BOOLEAN DEFAULT "false"
}
```

This creates a record named `mouse` containing five scalar nodes. For instance, `model` is a plain string variable, while `deviceFile` is a special string type whose definition is shown above the record. The first two lines contain some syntactic sugar that need not concern us at the moment.

%Those funky strings (|":MOUSE_EMULATE_NONE:")|
%are tags for the COAS message catalogs.

This definition would be stored in a file named `peripherals/mouse.schema` (usually below `/usr/lib/coas/schema`) so that the mouse configuration would be accessible by the name `peripherals.mouse.device`.

When accessing data items, COAS instantiates the portions of the instance tree from the schema definition and populates the data by invoking so-called "mappers". These mappers are responsible for parsing and writing back system files, locking them if necessary. Usually, they are written in C++ and kept in shared libraries loaded on demand. The most recent release also supports mappers written in Python.

In the case of the mouse device, there is no standard location where this information is stored. On a Red Hat box, for example, it is kept in `/etc/sysconfig/mouse`, a file which contains a list of shell variable assignments. COAS already has a general-purpose mapper for this type of file (it turns out that about 80% of all system files are quite close to four or five standard formats), so all that is left is

defining the mapping. This is done by the so-called platform repository, where we might enter code like this:

```
peripherals.mouse.device {
    mapper      builtin.sysconfig
    path        /etc/sysconfig/mouse
    relation     MOUSETYPE:model:\
        PROTO:protocol:\
        DEVICE:deviceFile:\
XEMU3:emulate3btn(map=/no=false,yes=true/)
}
```

The `mapper` keyword associates the mapper specified with the mouse device node. When accessing the device node for the first time, COAS detects this and invokes the mapper in order to populate the tree below the mouse device node. The mapper retrieves the `path` parameter and reads the file specified. The `relation` parameter tells the mapper which shell variables within the file correspond to which data model nodes.

The same thing happens when you have modified a protocol (e.g., the mouse) and invoke the device node's `commit` function. The data engine will invoke the mapper in order to write the data back to the file. Again, the mapper will use the specified `relation` to determine which data model values will be assigned to which shell variables. Note that in an act of vi-administrator friendliness, the mapper does not touch shell variables it does not know about and tries to preserve comments as well as it can.

The platform information is usually installed by merging it with the main COAS platform definition, which resides in `/usr/lib/coas/repository`.

User Interaction

Having written and installed the above files, you can already display and modify the mouse configuration using COAS. For example, COAS comes with small utilities such

Listing 1. Mouse Configuration Module

```
class Mouse(clam.CLAM):
    def __init__(self, ui, name):
        # Base initialization - UI handle,
        # dialog name, NLS name
        clam.CLAM.__init__(self, ui, name,
            "peripherals/mouse")
    def run(self):
        self.mouse = dm.InstanceLookup(
            "peripherals.mouse.device");
        # Create prompt dialog
        d = self.promptDialog("mouse");
        # Add edit fields:
        d.addInstancePrompt(self.mouse,
            "model<15>")
        d.addInstancePrompt(self.mouse,
            "protocol<15>")
        d.addInstancePrompt(self.mouse,
            "deviceFile<15>")
        # Get marker for change log
        marker = self.mouse.getMarker();
        done = 0;
        while not done:
            # Execute dialog
            result = d.execute();
            if result == "true":
                done = self.commit(self.mouse, marker);
            else
                done = self.cancel(self.mouse, marker);
```


as coas dump and coas change that let you dump portions of the data tree or modify individual nodes. You can even write Python scripts that perform more complex operations on your data.

However, the ultimate goal (at least for us) is a Python module that interacts with the user, guiding him through the administration task. The module sits on top of the database engine and operates exclusively on the abstract data representation. It displays data to the user, selects which items are to be edited, provides on-line help, etc.

Why Python? Well, a very early prototype used Tcl as the scripting language, but for various reasons it didn't work too well. In contrast to Tcl, Python has fairly good object support and at least as good an extension mechanism. The other candidate was Perl, but we decided against it because it is so easy to write horrible code in Perl.

Communication with the user happens via an abstract user interface API written in C++, which currently supports a curses and a Qt front-end. Work on an extended Qt front-end that takes advantage of features provided by KDE is in progress. Of course, in order to be able to use this API from Python, a Python wrapper is provided.

The user interface provides a limited but useful set of dialogs: notice/question dialogs consisting of a text and a few buttons; list dialogs (single- and multi-selection, with or without check boxes, etc); prompt dialogs (containing edit fields for one or more scalar values); and table dialogs (which display data in a table, allowing in-place editing).

For instance, a minimal module for editing the mouse configuration would look like Listing 1 (some of the Python fluff, such as import statements, is not shown). For those not familiar with Python, this code defines the class `Mouse`, derived from the `CLAM` class defined in module `clam`. The `{__init__}` method is Python's way of declaring a constructor.

The method `run` is invoked by COAS. The first thing it does is look up the data model node for the mouse device. As described above, this step will trigger the parsing of the configuration file into the internal data representation.

Next, a prompt dialog is created and three edit fields are added for

the mouse's model, protocol and device file. The last few lines are the somewhat standard dialog loop. Depending on whether the user terminates the dialog by pressing the Okay or Cancel button, either the `commit` or `cancel` method (inherited from the `CLAM` base class) is invoked, which displays a small question dialog along the lines of "Do you really want to save/cancel?"

What about Labels?

The first thing that will probably strike you as odd about this example is that it has no label strings anywhere. Nevertheless, the dialog is supposed to have a title, edit fields are supposed to have a label to their left, etc.

The answer is that COAS generates NLS strings for you out of the information it has. For instance, when creating the prompt dialog, we inconspicuously passed the string `mouse` into the function. As a consequence, COAS will create tags such as `|":MOUSE_TITLE:"|` for the dialog's title and attempt to look it up in the module's message catalog. (The message catalog name was specified in the class constructor.) Likewise, for the protocol edit field, it will gen-

erate the tag `|":MOUSE_PROTOCOL_LABEL:"|`. All you need to do is write the message catalog, mapping these funky strings to intelligible English (or French, German, etc.) and install the file.

Editing Data

Looking at the sample code above, you may also have thought: I understand where they put the data in the dialog, but how do they put it back into the data model?

This is the interesting part about the data editing process. If you have ever programmed Motif, you know how tedious it can be to extract the value to be edited from the data model, put it into the dialog and write the resulting value back to the data model when the user hits the OK button.

The approach taken by COAS is to tie data model nodes into the dialog directly and let the dialog select an appropriate widget type (string, combo box, toggle button, spin button, etc.). When the user provides a new value, the dialog will automatically check the value's syntax against data model constraints and write it back into the data model.

The world's fastest SCM system

runs on the world's best operating systems

Linux on Intel Linux on SPARC
Linux on Alpha MkLinux for Mac

(plus twenty other UNIXes, Windows, Mac, and more)

PERFORCE

THE FAST SOFTWARE CONFIGURATION MANAGEMENT SYSTEM

info@perforce.com www.perforce.com
2411 Santa Clara Ave., Suite 40 Alameda, CA 94501 1-510-864-7400

© 1998 by Perforce Software, Inc.

Dependency Model

When operating on system data, constraints and/or dependencies present are:

- When enabling NIS, nis client must be enabled.
- Package XYZ requires user.
- When configuring PPP, select from list of configured modems.

Some of these dependencies can be coded into the CLAM; some cannot. COAS supports different mechanisms by which these can be implemented:

- Calling other CLAMs, e.g.,
`call hardware.modem.selectModem()`.
- Action routines: these action routines can be tied to individual data nodes and are invoked when committing data back to system files.
- Pre-commit verify (e.g., really delete this user *abc*?)
- Pre-commit actions (e.g., take removed IP interface down now?)
- Post-commit actions (e.g., populate home directory of new user)

In our example, the dialog would create a simple string edit field for `deviceName`, a pop-up list for `protocol` (since it is limited to a set of choices) and a toggle button for emulation.

What's more, this mechanism offers you easy-to-use context help for each input field, bound to the F2 key. Adding this type of help to a data item is as easy as adding the `HELP` attribute to the data definition in the schema file:

```
device      RECORD {
    model     STRING HELP "HELP_MODEL"
    protocol  MouseProtocol HELP "HELP_PROTOCOL"
    ...
}
```

These help messages will be looked up in the message catalog associated with the schema file (remember the `MSGCATALOG` keyword in the schema file?) and displayed in a pop-up dialog whenever the user presses F2.

Of course, every scheme you devise has a drawback. In this case, it is how to cancel changes made during the execution of the dialog. When the user presses the Cancel button, he wants all changes to go away.

This is where the marker object comes into play. The data node's `getMarker` method obtains a marker for the node's change log (called a journal in COAS lingo). When the user requests a discard of all changes, the CLAM base class invokes `self.mouse.cancel(marker)`, which reverts all changes made after the marker object was obtained.

Where's the Beer? er, Beef?

I have to admit that the above example, in its simplicity, is a bit deceptive. What I'm showing here is the simplest version of a dialog. In fact, what you see here is just a glorious interface to the configuration file because it does not offer the user any help or guidance. A good dialog would automatically choose the appropriate device file when a selection is made (e.g., a bus mouse) and keep the user from enabling three-button emulation for mice that already have three buttons. As a consequence, your average COAS module will have a lot more than those 20-odd lines in the example above.


However, the greatest advantage COAS offers in this context is that it relieves you of the usual hassle when working with a GUI and lets you concentrate on the data flow instead.

Are You Curious?


If this article has piqued your interest and you would like to take a closer look at COAS, you can find out more about it on <http://www.coas.org/> and <http://developer.coas.org/>. If you want to participate in the development of COAS, don't hesitate to contact me. ■



Olaf Kirch has been a Linux enthusiast since the MCC Interim days and has authored the *Linux Network Administrator's Guide* as well as various pieces of software for Linux. He has been the principal maintainer of the Linux NFS code for several years and has been working for Caldera since 1997. He can be contacted at okir@caldera.de.



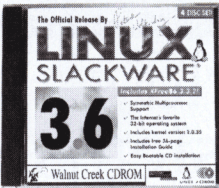
Walnut Creek CDROM



Linux Slackware® 3.6 (4 disc set) \$39.95


Slackware® Subscription — Save!\$24.95

We will automatically ship you updates. You may cancel at any time.




Unlock the power of your PC with Linux Slackware® 3.6!

- Includes XFree86 3.3.X
- The latest kernel
- Free 36-page installation guide!




Linux™ Toolkit (6 disc set):

Get the tools to make your Linux operating system complete! You get 3.4+ gigabytes of apps, source code, and development tools. ...\$29.95












Slackware® Installation Manual:

This 345 page book and current Slackware release disc contains everything you need to get started with Slackware Linux. ..\$29.95



*Linux: The Complete Reference: NEW EDITION!

This 1,608 page book is the most complete and current collection of info on Linux available anywhere!\$39.95



Order Now! 1-800-786-9907

Walnut Creek CDROM

\$5 shipping per order (\$9 overseas)

Phone: +1-925-674-0783 Fax: +1-925-674-0821
Email: info@cdrom.com URL: <http://www.cdrom.com/>
4041 Pike Lane, Suite F-475, Concord, CA 94520, USA

Many more 1 year *unconditionally* guaranteed CDROMs available!
Please call or email today to ask for your *free* catalog.

475


LINUXWORLDTM

CONFERENCE & EXPO

KEYNOTE PRESENTATIONS:

Dr. Michael Cowpland, President & CEO, Corel Corporation
Mark Jarvis, Senior Vice President of Worldwide Marketing, Oracle
Linus Torvalds, Creator of Linux

On Top of the Wave



Introducing LinuxWorld

Conference: March 1-4, 1999

Expo: March 2-4, 1999

**San Jose Convention Center
San Jose, CA**

Now Linux has a brand-new show of its own. LinuxWorld Conference & Expo makes its debut March 1, 1999. It's going to be a blockbuster, because the Linux OS is the most talked-about technology in the IT industry. Attend LinuxWorld at the San Jose Convention Center and prepare to step into your own starring role in this rapidly emerging industry!

Geared to developers, enthusiasts and professional end-users, the Conference program will feature issues and solutions focused on Linux technology. There will be fifteen pre-conference tutorials and over 40 conference sessions. Attendees will include IS managers and decision-makers from Fortune 1000 companies, Linux developers, ISVs, VARs, OEMs and system administrators. The Expo will feature more than 100 emerging companies and organizations dedicated to Linux products and services.

Reserve your seat now! For more information, complete and fax or mail the coupon or visit:

www.linuxworldexpo.com

On the Web site, you'll get the latest updates, including conference session descriptions, show highlights, the exhibitor list, and more.

**YES! Tell me more about
LinuxWorld Conference & Expo '99**

in San Jose. I'm interested in

☐ **Attending** ☐ **Exhibiting**

Name

Title

Company/Organization

Address

City/State/Zip

Phone Fax

email

Mail to:
LinuxWorld Conference & Expo
P.O. Box 339
West Bridgewater, MA 02379
FAX: 508.620.6668
or Call 800.657.1474

Owned and Managed by:

IDG
WORLD EXPO

PCW

Flagship Sponsor:

Linux
WORLD

Flagship Sponsor:

LINUX
JOURNAL

Media Sponsor:


EARTHWEB

Media Sponsor:


developer.com

Media Sponsor:

INFO
WORLD

Media Sponsor:

Macworld

Media Sponsor:

NetworkWorld

Media Sponsor:

Sys Admin

Csound for Linux

Mr. Phillips discusses some history as well as what's happening now in the Linux Csound world.

by David Phillips

Csound is a music composition and sound programming language originally written by Barry Vercoe at MIT. As Nick Bailey pointed out in his October 1998 *LJ* article "Sculptor: A Real Time Phase Vocoder", Barry's original MUSIC11 program was eventually ported from PDP-11 assembler to UNIX C, where it became Csound. MUSIC11 was derived from the pioneering MusicV program by Max Mathews, perhaps the most revered "Founding Father" of computer music technology.

One of MusicV's major innovations was the implementation of the unit generator, a "black box" concept that allowed great extensibility to the language. A unit generator can be a signal generator or modifier, a patching opcode, a sensor, or it can provide sound file I/O and signal display types. Csound has evolved into a notable successor to Music V, quickly accommodating new synthesis methods and DSP algorithms. It is now at the cutting edge of modern computer music

software. Linux Csound has done more than simply kept pace with that evolution—it offers capabilities not found with versions available on other platforms.

Enter Linux

In 1996, I wanted to try out the Linux OS. I knew certain software synthesis languages would run under it, and those languages were not available for DOS/Windows machines. Although Csound does indeed run under Microsoft operating systems (and many others), I was interested in seeing how well it would run under Linux. Jonathan Mohr had already added the real-time audio support for Linux, but I immediately discovered I had stumbled upon another big "DIY" (do it yourself) project. The source code available from the Bath, UK FTP site (the primary repository for the "canonical" packages) was a general UNIX package, without Linux-specific Makefiles or any other compilation amenities. Although I was a novice at both Linux and the C programming language, I jumped in and started thrashing. With good assistance from John Fitch (maintainer of the Bath site and the canonical sources) and the helpful members of the Csound mail list, I finally produced a working set of Makefiles for the entire source tree. I soon had a fast Linux Csound with full support for X displays, real-time audio output and all the current opcodes. Professor Burton Beerman kindly provided an FTP site for my Linux Csound packages on his MusTec server at Bowling Green State University, and for two years I maintained the public version on that site and at Bath.

Linux Csound: the Plot Thickens

Early in 1998, I received a message from Professor Nicola Bernardini at AIMI (Associazione di Informatica Musicale Italiana). He had thoroughly rewritten the Linux Csound Makefiles and wondered if I might be interested in adding them to the source package. His offer came at a good time, as I knew the code maintenance needed a more solid structure. Nicola's expertise was just the right factor appearing at just the right moment. His Makefiles enabled me to quickly prepare a variety of distribution packages (with or without X support, static build or shared lib, real-time audio enabled/disabled, etc.) and compile a more complete build of the source tree. Most importantly, the Makefiles created libcsound.so, a shared library which drastically reduced the binary's memory footprint (from about 450KB to less than 20KB).

Around the same time, developer Gabriel Maldonado wrote a set of MIDI output opcodes, allowing Csound to be used as a MIDI composition/control instrument. Csound already accommodated MIDI input, directly from /dev/midi or from a Type 0 Standard MIDI File (see "Real-time MIDI Input"). Gabriel's opcodes are different: they permit exploration into MIDI composition algorithms simultaneously with the rest of Csound's real-time I/O. Hypothetically, it would be possible to have a MIDI device

Each of these applications was built using freely available tools.

Csound in a Nutshell

Csound is sometimes referred to as a waveform compiler, so it requires two types of source files, orc and sco. An orc file designs a Csound instrument, which might be anything from an FM oscillator (foscil) to a sample playback engine (loscil or soundin). The orc file may internally hard-wire its necessary parameter values, or it may receive them from the sco file. An sco file specifies event start times and durations and any required stored-function tables. The finished orc/sco files are then compiled by the **csound** command. In the following example, they are compiled to a WAV format sound file:

```
csound -o foo.wav -W foo.orc foo.sco
```

When running Csound in an X session, a small window will pop up with a graphic display of the function tables in the sco file. Clicking over the window will allow compilation to begin and a series of messages will appear in the starting xterm. If no errors are reported, a new sound file was created and it is now available for further processing in Csound or one of the many Csound utilities.

By the way, beginners are often surprised to learn that despite its name, Csound presents itself in an archaic syntax more like assembler or FORTRAN than C. The "C" part of the name refers to the fact that the Csound source code is written in C.

Real-time Linux Csound

A simple introduction to Linux Csound's real-time usage can be presented by defining an instrument (a simple oscillator), creating it in a text editor and saving it as `my.orc`:

```
sr=22050 ; sets sampling rate
kr=441 ; sets control rate for k-rate components
ksmps=50 ; sr/kr, number of samples per control
          ; period
nchnls=1 ; monaural output
```

```
instr 1
kamp = 10000 ; raw amplitude
kfreq = 440 ; an A at 440 cps
ifn = 1 ; stored function table 1
asig oscil kamp, kfreq, ifn ; an audio signal is
; created by an oscillator playing
; a stored sine wave at kamp and kfreq out asig
; the audio signal is sent out endin
```

Next, we write another file called `my.sco`. When compiled, the score will be played by the instrument:

```
f1 0 8192 10 1 ; stored sine wave
i1 0 3 ; instrument 1 plays for 3 seconds,
          ; start time
at 0
e
```

We then compile our `orc/sco` files into a sound file:

```
csound -omy.wav -W my.orc my.sco
; creates a WAV format sound file
```

Csound will create the sound file, the user can play it with `vplay` or any other WAV player and it can be edited by `MiXViews`, `DAP`, `Snd` or any other Linux sound file editor.

If we wish to send the sound output directly to the sound card DAC, we use this command:

```
csound -o devaudio -W -dm6 my.orc my.sco
```

where:


- `-o devaudio` sends output to audio device in real time
- `-W` creates WAV format (uses `/dev/dsp` for audio device)
- `-dm6` turns down messaging for better real-time response

With a full-duplex sound card, it is possible to have simultaneous audio input and output, allowing real-time signal processing. A command for such a setup (using the ALSA driver) would look like this:

```
csound -iADC -oDAC -W -dm6 inout.orc inout.sco
```

where:

- `-i ADC` input to sound card analog-to-digital converters
- `-o DAC` output to sound card digital-to-analog converters



YOUR LEADING SOURCE FOR COMPUTING AND NETWORKING NEEDS www.asacomputers.com Since 1989

NETWORKING PRODUCTS

- CISCO Premier Reseller
- A Variety of Networking Solutions
- Other manufacturers we feature:
3COM, Ascend, Bay, Cybex, Cyclades, Digi, Livingston

RAID SOLUTIONS

- Specific for different OS's.
- PCI/SCSI and SCSI/SCSI Options
- Adaptec, CMD, DPT, Mylex, AMI

- Start up Kit for ISPs
- X-Windows workstations for development
- Firewalls and LAN/WAN Gateways
- RAID Solutions
- Rackmount and Stand-alone Configurations
- Laptops preloaded with specified O.S.
- Redhat/Debian/Slackware/S.u.S.E/Caldera/Linux
- Also: BSDI and FreeBSD/NT/Solaris.



ASA COMPUTERS, INC.

200 East Trimble Road, San Jose, California, 95131 USA

TEL: 408-232-5999 X 201 • FAX: 408-232-5959

E-mail: linux@asacomputers.com • URL: www.asacomputers.com



controlling one Csound instrument while another instrument sends its output to devaudio. Given support for a full-duplex sound card, it should even be possible to have asynchronous I/O for both the MIDI and the audio ports.

Alas, no routines had been written for Linux Csound that would accept the data from Gabriel's opcodes and send it out to the MIDI port. After studying John Fitch's code for the Windows Csound MIDI output handler, I decided to try writing the appropriate calls for Linux. I fumbled around with the OSS/Free API and eventually wrote the code needed to activate the requested MIDI interface and accept the control data sent to it from the Maldonado opcodes. Linux Csound was as up-to-date as any other version, and the necessary code for MIDI output had been trivial to write, consisting primarily of a few calls to the sound card API macros.

The CVS Repository

The next major step taken for Linux Csound was the establishment of a CVS repository. I had been complaining to Nicola that I found myself constantly checking everything coming to me in the canonical UNIX package, when he suggested the need for a revision control system. He volunteered to set it up at AIMI and after some trial-and-error hacking, he established the system we work with today. The

CVS repository maintains separate directories for the canonical sources and the Linux-specific code. In this way, we can avoid rewriting sources just for Linux and we are always able to refer back to the "untouched" originals. Anonymous access to the CVS is permitted, but submissions for changes are carefully screened by the maintainer.

The Csound UNIX/Linux Development Group

Of course, a CVS development repository isn't of much use unless it has developers contributing, so a logical next step was the formation of the Csound UNIX Development group. Programmers Robin Whittle and Damien Miller joined in immediately, and Damien kindly provided a web page with all pertinent information for anyone interested in joining the group. It is worth noting that the group is for development, not just developers. We welcome anyone interested in seeing Linux Csound grow into the finest language of its kind. Programmers are certainly welcome, but so are musicians, audiophiles, DSP engineers and anyone else with an interest in Csound and its possibilities.

In October 1998, two new members made significant contributions to the group's activities. Gabriel Maldonado donated his entire source tree to the CVS repository, which enables Linux Csound to keep up with the developments for his Windows versions. (This generosity is quite typical of the

Real-time MIDI Input

Now, assuming you have a MIDI interface properly installed, you can play your first instrument by changing the sco file like this:

```
f1 0 8192 10 1
f0 10000 ; dummy table, keeps the instrument
          ; active for 10000 seconds
e
```

and compiling your orc/sco files so:

```
csound -o devaudio -W -dm6 -M /dev/midi my.orc
my.sco
```

where -M /dev/midi tells Csound to look for MIDI input instead of score sound events.

Another way of triggering my.orc is to use a Standard MIDI File. The orc/sco can be used as in the last example, but the command line changes to:

```
csound -o devaudio -W -dm6 -F my.mid my.orc my.sco
```

where -F my.mid tells Csound to expect input from the MIDI file.

If we want to make our sounds more interesting, we can map MIDI values into the orc file and arbitrarily utilize that data as shown in Listing 1.

Our score file (sco) is as before, but with an added function table:

```
f1 0 8192 10 1 ; sine wave
f2 0 8192 10 1 .9 .8 .7 .6 .5 .4 .3 .2 .1
; ramp wave
f0 10000
e
```

Once again, these files can be compiled for control by either a MIDI input device or a MIDI file. They demonstrate that simple Csound components can easily be

built into elaborate sounds and control structures. I have used the simplest instrument here, but already it is quickly evolving toward greater complexity.

Real-time MIDI Output

Using Gabriel Maldonado's MIDI output opcodes is a little different. The following example demonstrates a more complex control system as shown in Listing 2. The sco file for this instrument is as follows:

```
; MIDI OUT opcodes
; mussle.sco

t 0 60 ; tempo indicator
i1 1.0 1.0 ; the different p3 durations will
; stretch out
i1 2.0 7.0 ; the play of the notes indicated by
; knum
i1 9.0 18.0
i1 18.0 17.0
i11 0.0 1
i100 0 1
e
```

The command line for compiling these files looks like this:

```
csound -Q0 -n -dm6 mussle.orc mussle.sco
```

where -Q0 selects the first MIDI device and -n suppresses write to disk (for performance enhancement).

Given a fast enough CPU and disk, it is possible to combine the various real-time I/O options. I have tested simultaneous real-time audio with MIDI output opcodes; it works, but with a big performance hit. Considering that my machine is a lowly AMD 486/120, it is to be hoped that a fast Pentium would lessen that hit.

Listing 1. Mapping MIDI values—orc File

```
instr 1
inum notnum ; note number
kfreq cpsmidib ; MIDI to frequency
iamp ampmidi inum*100 ; MIDI to amplitude
                        ; (scaled within range)
if inum > 60 goto fun1 ; if the MIDI note number
                        ; is over 60...
if inum < 61 goto fun2 ; if the note number is
                        ; less than 61...

fun1:
ifn = 1
goto contin
fun2:
ifn = 2
goto contin
contin:
asig oscil iamp,kfreq,ifn
out asig
endin
```

Csound community. Much code sharing occurs on the Csound mail list, with new instrument designs freely offered, along with much healthy debate over various computer music issues.) The other signal addition has been Fred Floberg, whose contributions require special description.

Csound's internal support for real-time audio has been dependent on calls to the API for the OSS sound-card drivers. While certainly sufficient for casual use, many sonic notions such as full-duplex and multiplexed real-time audio I/O are not realizable by the OSS/Free driver. However, the ALSA driver does indeed support those uses; thanks to code from Fred Floberg, Linux Csound now explicitly supports the ALSA interface. (The ALSA project, led by Jaroslav Kysela, is forming a new extended sound system API compatible with OSS/Free, but permitting much more advanced uses for sound-card features not supported by OSS/Free.) Fred is currently working on expanding MIDI file support. Csound now supports only Type 0 MIDI files, but Linux Csound should soon support the Type 1 and perhaps even the Type 2 Standard MIDI File formats.

Also, thanks to Robin and Damien, the Linux Csound distribution now supports the popular RPM packaging and can be built for glibc (libc6) systems. Debian users will be pleased to note that developer Günter Geiger has prepared a package in the DEB format. Finally, Nicola Bernardini has written a Csound orchestra (instrument design) parser, which we hope will eventually be absorbed into the package. Such a utility is most helpful to a GUI designer, which brings me to my next topic: the power of Linux Csound and X.

The X Picture

My Linux soundapps web page shows more than twenty entries in the "Csound Helpers" section. The brief descriptions which follow are just that—brief descriptions which in no way indicate the full power of these applications. The examples shown here are for Linux systems running X; some excellent command-line utilities exist too and are included on the Linux soundapps page for those tools. All of these utilities work with the current versions of Linux Csound (3.47 or higher).

Note that each of these applications was built using freely available tools. The GNU C and C++ compilers, Tcl/Tk, Java, LessTif and WINE are powerful allies in the advance-

Listing 2. MIDI Output Opcodes—orc File

```
; MIDI OUT opcodes
; mussle.orc
sr = 44100 ; also experiment with other sr and
            ; kr rates
kr = 44100
ksmps = 1
nchnls = 1
; This instrument shows how the moscil opcode
; can be used for a sort of
; "indeterminacy" aspect in composition. knum
; and kvel control the output
; of MIDI notes and their velocities over p3
; (note-stream duration as specified
; in the score). kran output is used for kdur
; (actual duration) and kpause
; (pause between notes).
;
; The moscil (MIDI oscillator) opcode permits
; control of MIDI channel, note-number,
; velocity, duration, and time between notes.
```

```
instr 1
kchan = 0
knum linseg
40,p3/7,60,p3/7,46,p3/7,54,p3/7,24,p3/7,58,p3/
7,48,p3/7,80
kvel linseg 20,p3/2,39,p3/2,33
kran pcauchy 23
kdur = int(kran)+.46
kpause = frac(kran)
moscil kchan, knum, kvel, kdur, kpause
endin
```

```
; The next two instruments control Program
; Change messages sent to my external synth
; and my MIDI-controlled mixer.
```

```
instr 11
kchan = 0
kprog = 2.
kmin = 1.
kmax = 128.
koutpc kchan, kprog, kmin, kmax
endin
```

```
instr 100
kchan = 8
kprog = 6.
kmin = 1.
kmax = 128.
koutpc kchan, kprog, kmin, kmax
endin
```

ment of Linux sound and MIDI software. Their developers are to be commended for the wonderful work they have done for the good of the Linux community.

Cecilia

Cecilia (by Jean Piche and Alexandre Burton at the University of Montreal) is a fully-developed Csound composition and sound-processing environment. Written entirely in Tcl/Tk, Cecilia utilizes the entire range of possibilities afforded by Linux Csound, presenting a beautiful graphic interface (customizable, of course) and a powerful

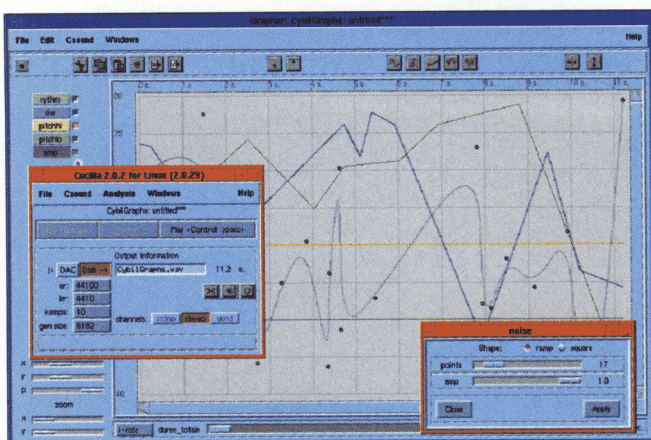


Figure 1. Cecilia

composition language (Cybil). Numerous real-time controls are supported, nearly all aspects of the program are user-definable, excellent on-line help is available and the GUI fully exploits Tk in the X environment. Cecilia won first place in the awards for computer-aided composition and realization software at the 1997 Second International Music Software Competition in Bourges. (See Figure 1.)

Rain

At the other end of the scale is developer Matti Koskinen's **rain**, a GIF-to-Csound score converter. A Csound score is the control file for a Csound instrument, providing it with such values as event start times, durations, amplitudes and frequencies, waveform selection and so forth. Matti's utility simply takes a GIF image, applies some user-defined values and magically translates it into a Csound score. The score can then be synthesized and played from within the application, or it can be saved to disk for later processing (perhaps in Cecilia). (See Figure 2.)

Adsyn

Adsyn is a graphic editor for Csound "hetro" analysis data files. **hetro** is one of the Csound sound file utility programs and its operation is quite simple. Using a heterodyne filter bank, it analyzes a sound file and creates a data file of separated frequency and amplitude values. That data file can be read and graphically represented by Adsyn and the frequency and amplitude components can be freely altered using the mouse. Csound's resynthesis opcode (adsyn) can be called; the edited file can then be synthesized and played from within Adsyn. Professor Oyvind Hammer originally wrote Adsyn for SGI machines at NoTAM, a Norwegian center for music and acoustics research. With his good graces, I began the port to Linux. It was finished with much assistance from Nicola Bernardini. (See Figure 3.)

Ceres2

Ceres2 is Johnathan Lee's enhanced version of Oyvind Hammer's Ceres, described in my September 1998 *LJ* article "Porting SGI Audio Applications to Linux". Johnathan greatly extended the editing capabilities of the original software engine, which essentially performs a Fast Fourier Transform (FFT) on a sound file and renders a graphic representation of its frequency content and activity. The graphic display can be edited in various ways, a large number of transforms (spectral mutations) are available, up to three graphic linear control functions may be specified and a variety of output formats are supported, including two

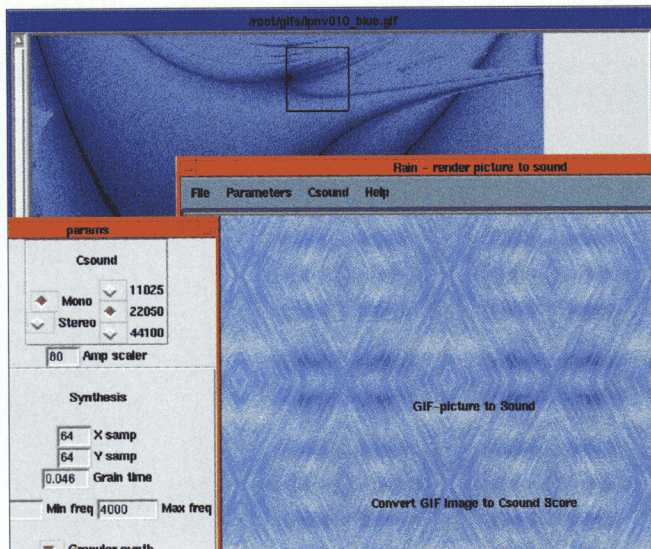


Figure 2. Rain

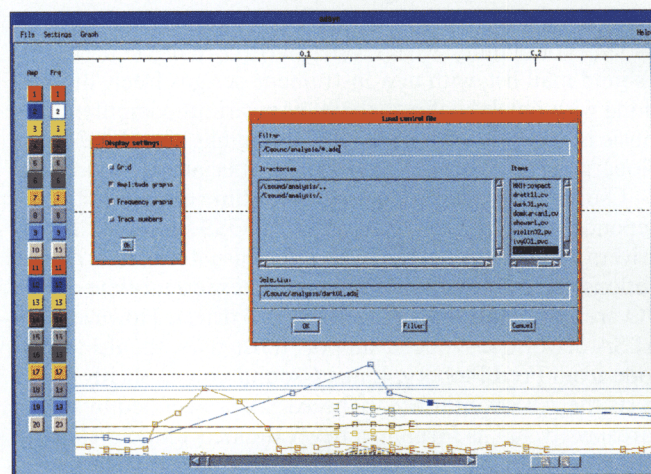


Figure 3. Adsyn

types of Csound scores. Ceres2 also extends some of the command-line analysis variables such as FFT size, analysis window size and window overlap. The Linux port was done by me, but it was dependent on work already done on the original Ceres with great help from Richard Kent, who also supplied the invaluable tichstuff libraries which replace the SGI libs. (See Figure 4.)

Rosegarden

The Rosegarden suite includes a MIDI sequencer, a common-practice music notation display and the very nice feature of being able to save your work as a Csound score file. Such a tool is especially helpful for users most comfortable with standard notation conventions, allowing them to compose with their familiar symbols and then easily convert their creations for use with Csound instruments. (See Figure 5.)

HPKComposer

The Java programming language lends itself to the easy creation of platform-neutral user interfaces. Jean-Pierre Lemoine's HPKComposer is an excellent example of a "pure Java" application, running under Windows, Mac OS and UNIX variants. Preparation for Linux is straightforward, depending upon successful installation of the Java development environment (JDK) in version 1.1.6 or higher, the Swing class libraries (version 1.1 beta3) from Sun

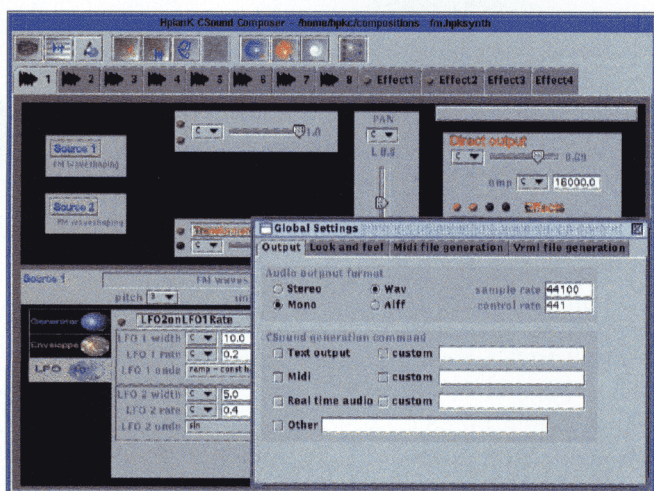
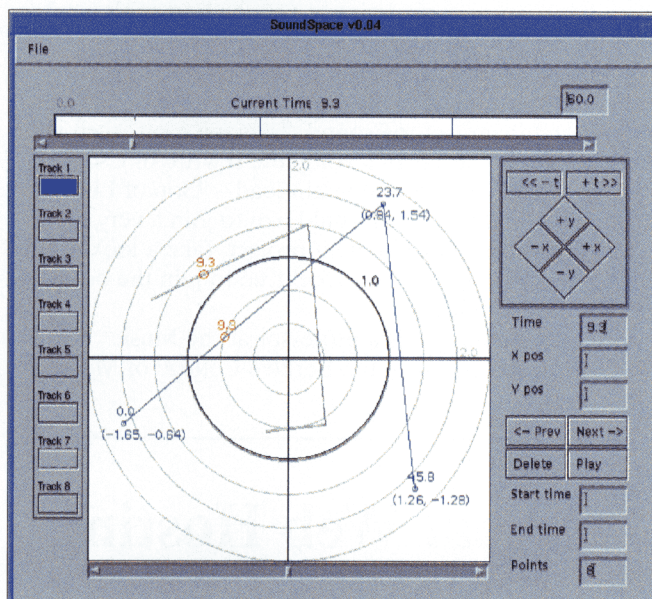
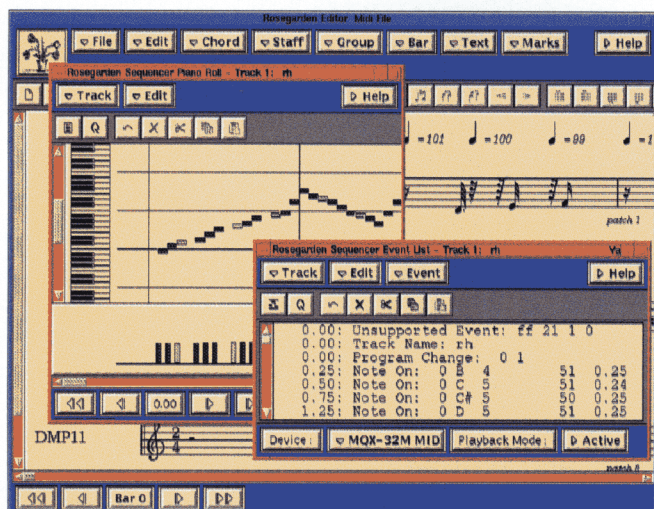
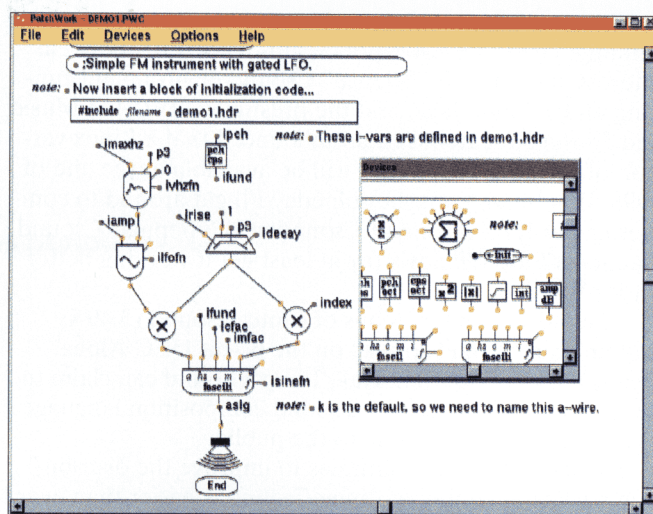
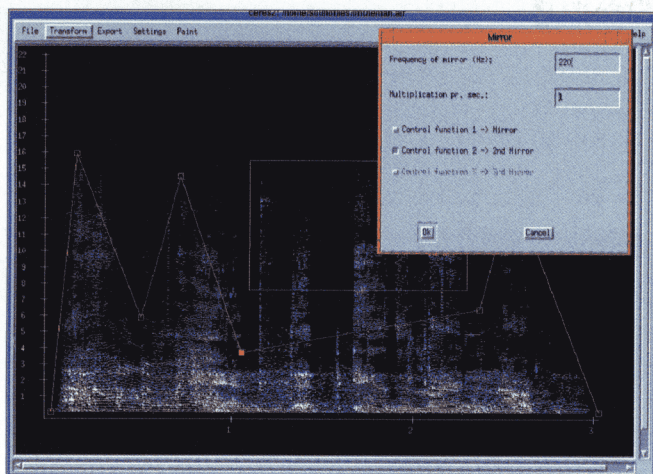


Figure 6. HPKComposer

Microsystems and Csound. HPKComposer blends aspects of the CMask program with the synthesis and DSP methods of Csound: tendency masks are used to create composition algorithms, which are realized by the synthesis engines (opcodes) of Csound. VRML displays are supported, the program is user-extensible, and although Java's current sound support is limited to 8-bit 8 kHz audio, when JDK 1.2 arrives it will support 16-bit 44.1 kHz CD-quality sound. (See Figure 6.)

PatchWork

Russell Pinkston's PatchWork for Win95 is a graphic "patcher" for the design of Csound instruments. Although a UNIX/Linux version of this program exists (XPatchWork), it has not been maintained and is in need of some serious debugging. However, the Linux WINE Windows emulator can run the Win95 version, proving once again that Linux always finds a way. (See Figure 7.)

SoundSpace

Developer Richard Karpen has generously shared many of his opcodes with the general Csound community, one of which is called “space”. In the Csound manual entry for space is a mention of a GUI for creating the values needed by the GEN28 stored-function table, and SoundSpace is that GUI. Written in core Java, this unique utility provides a visual interface for determining the placement and sonic trajectories of up to 8 sound files in the auditory space, with support for stereo and 4-channel output. (See Figure 8.)

Into the Future

What is still to come? By the time this article is published, I hope to have some more Csound/Java applications

running. Developer Michael Gogins has expressed great interest in seeing his "Silence" Csound environment running under Linux Java, and the prestigious IRCAM Music and Sound Research Center announced that a Linux version of their MAX for Java will be available at the end of 1998. Who knows; maybe someday I'll get around to completing my Tcl/Tk clone of Csounder, the popular Csound "launcher" for Windows (or at least get it working better under WINE).

The most recent versions of Linux Csound (3.49.xx and up) can be built for use on the 64-bit DEC Alpha. Thanks to developer Ed Hall, Linux Csound can claim to be the first 64-bit music and sound composition language widely and freely available to the public.

Nicola Bernardini continues to improve the distribution packaging: building Linux Csound is easier than ever, thanks to his incorporation of the **configure** utility. Work proceeds on accommodating **autoconf** and **automake**, since it is a primary objective to use the best tools available for creating the best possible distribution.

One of the intriguing problems facing the development group is how to make Csound re-entrant, enabling a plug-in architecture for Csound. To many of us, such an undertaking would mean a complete rewrite of Csound, and who knows where that might lead—"Son of Linux Csound", perhaps? If you would like to join a very interesting distributed development project, take a look at the links listed in Resources and feel free to join the development group mail lists.

Richard Boulanger is a professor at the Music Synthesis Department of the Berklee College of Music. In


the spring of 1999, his Csound book will at last be published by MIT Press. On one of the included CDs, you will find an article (which will, of course, be out of date by then) about running Csound under Linux. Yes, it was written by me, but I don't mention it to blow my own horn. This book is a massive tome and it includes contributions from all the major (and some not-so-major) members of the international Csound community. It should inspire many new users, several of whom will discover for the first time that Csound is available on the Linux platform.

Final Words

Linux Csound offers terrific possibilities for real-time computer music performance. Along with advances in real-time support, Linux Csound can be expected to stay at the cutting edge of synthesis methodologies, interface design, DSP algorithms and composition strategies. It is an ideal tool for contemporary sonic exploration and it demonstrates once again the flexibility and power of Linux, the cutting edge OS for the modern musician. ■

David Phillips (dlphilp@bright.net) is a composer/performer living in Ohio. Recent computer-music activities include an ambient composition for the artist Phil Sugden, lecturing on computer-music programming languages at Bowling Green State University, and maintaining the "official" version of Csound for Linux. Dave also enjoys reading Latin poetry, practicing t'ai-chi-ch'uan, and any time spent with his lovely partner Ivy Maria.

Virtual Web Hosting



HURRICANE ELECTRIC
INTERNET SERVICES

Your Industrial Internet.

\$24.95 /mo.

| | | |
|--|--|---|
| 50 MB <small>storage</small> | 3 GB <small>transfer</small> | 21 POP3 <small>emailboxes</small> |
|--|--|---|

Other plans available, including colocation & dedicated servers.

- Ⓜ True virtual hosting with a static IP
- Ⓜ Multiple OC-3 connections (155 Mbps)
- Ⓜ UPS battery backup and an emergency generator
- Ⓜ Staffed NOC 24 / 7
- Ⓜ Access via FTP, telnet, and ssh
- Ⓜ Secure web directory (SSL) and cgi-bin

408 282.1540

w w w . h e . n e t

Resources

Nicola Bernardini's Csound site at AIMI. The most recent versions should always be available here and at the MusTec site: http://AIMI.dist.unige.it/AIMICSOUND/AIMICSOUND_home.html.

FTP server for the Music Technology Department at Bowling Green State University. Many Linux sound applications are available from this site: <ftp://mustec.bgsu.edu/pub/linux/>.

Not just for Linux Csounders, my own page includes pointers to orc/sco collections, on-line tutorials, web interfaces, sites with RealAudio clips and other good stuff: http://www.bright.net/~dlphilp/linux_csound.html.

Home of the canonical sources, maintained by John Fitch: <ftp://ftp.maths.bath.ac.uk/pub/dream/>.

My collection of links to everything MIDI and audio for Linux: http://www.bright.net/~dlphilp/linux_soundapps.html.

Information on joining the Csound UNIX/Linux Development mail-list, also information regarding the anonymous CVS: <http://www.ilogic.com.au/csound/>.

Information page regarding the Linux Audio Development mail-list: <http://www.bright.net/~dlphilp/lad.html>.

We're Stepping Up!

There are good reasons SuSE Linux is Europe's #1 Linux distribution:

- Experience
- Quality Engineering
- Commitment to Development

***It's time to see for yourself the difference
German engineering can make!***

New for SuSE Linux 6.0:

- Star Office 5.0 - Personal Edition
 - glibc-based
 - egcs 1.1 for C++
 - gcc 2.7.2.1 for C
 - XFree86™ 3.3.3, with support for Matrox G200 and G100 graphic adapters
- ...and much more!**

Contact us today!

*For inquiries within the USA
and elsewhere outside Europe:*

SuSE, Inc.
458 Santa Clara Avenue
Oakland, California 94610
USA
Tel +1-888-875-4689
+1-510-835-7873
Fax +1-510-835-7875
www.suse.com

For inquiries within Europe:

SuSE GmbH
Schanzäcker Strasse 10
D-90443 Nürnberg
Germany
Tel +49-911-7405331
Fax +49-911-7417755
www.suse.de



SuSE

Linux 6.0

Your complete resource for deploying the Linux operating system!

Software Libre and Commercial Viability

Mr. Rubini gives us his opinion of the Open Source movement.

by Alessandro Rubini

Fortunately, Linus' project of world domination is going to come true fairly soon. The trend toward this goal can be verified by checking how the press is behaving towards GNU/Linux solutions, looking at how several educational entities are going to introduce free software in the schools and verifying Linux's usual technical excellence.

Today in 1998 (yes, it is still 1998 as I write), the most important job remaining, in my opinion, is propagating the social and commercial implications of free software. While I greatly appreciated Russell Nelson's article "Open Source Software Model" in the July issue of *LJ*, I feel the need to expand on the points he briefly touched.

Please note that I'm not an expert in economics or politics. I'm just a build-it-yourself kind of technical guy who is reporting his own experience in the battle for survival, in the hopes of helping someone else adapt to new environmental conditions. Some of these ideas have already been discussed with friends or on the Free Software Business mailing list (fsb-subscribe@crynwr.com), which I joined after reading Russell's article.

Viability for Individual Consultants

The best feature of any computer system is flexibility—allowing users to tailor its behaviour to their own needs. This flexibility is often completely unknown to the general computer user, because proprietary software solutions tend to hide functionality behind a rigid external interface which denies any divergence from the expected behaviour—a *user's* behaviour.

When adopting free software, users are able to discover the real power of computer systems. Today I talked with a commercial consultant who never thought that programs could be adapted to one's needs. He confessed his company has always acted the other way around—they adapted their needs to the software they use. Most users are victims of their software and don't even realize it.

Educating the user base about the extendibility of software will open new markets to independent consultants, creating new employment opportunities. Every user has different needs and solving these needs often means calling for technical support from people who tailor or enhance the relevant software. While this is not even imaginable with proprietary programs, source availability allows any problem that might arise to be quickly solved and new features to be easily added. While you may think this would quickly lead to a *perfect* software package, individual needs are so diverse and specialized that the *perfect* package will never exist.

For example, I and others wrote a program for a local physiology center to analyze data for a typical kind of experiment. During two years of use, the physicians found so many ways to enhance the program that it is now reported as better than the commercial solutions. The total of all fees they paid during these years reveals the program to be more expensive in the end than some of the commercial alternatives. This fact is not relevant to my clients, as they have exactly what they want and they know they can have more should the need arise. The program is obviously GPL and other centers expressed interest in getting a copy.

As more and more people are choosing free software to address their needs, I'm sure some software companies will try to demonize Linux and the open-source movement because they are losing market share. Such companies will probably try to demonstrate that IT employment is decreasing and that humankind is being damaged by the general adoption of free software. This whole argument is bogus; computers exist to be programmed, and the more you allow programming them, the more you build employment opportunities. If you count the number of people who offer free software consulting, you will greatly exceed any shrinkage of proprietary companies. Sticking to my previous example, the physiology lab hired my company to write the program, and other centers interested in the product are willing to hire a local consultant for installing, maintaining and enhancing our package. Did I say "enhance"? Isn't the program working? Yes, the program is working well, but there *is* room for enhancement of the product. The local lab decided to stop development "because we must run our experiment rather than invent new software features". As anyone knows, every program has a bug and a missing feature, and this is where we build our credibility—bugs *can* be fixed and features *can* be implemented. As I suggested before, the more you make things programmable, the more they will be programmed.

Why should there be more employment opportunities in IT than there are now? First of all, because free software users have more requests for new features than users of proprietary products do, as explained above. Next, because anyone can build her own professionalism without paying to access the sources of information. I built my Linux expertise by studying source code and trying things out on my own low-end PC. Now I am confident I can solve any problem my clients might have, and my clients know I can (provided I am given enough time to deal with the problem).

Another critical point in addition to source availability is standardization on file formats, a field where proprietary products are revealing their worst features. Let's imagine

an environment where every file format in the system was known: you could, for example, create indexes from any document that is produced, thus easing later retrieval. This can be accomplished off-line without any load on non-technical personnel.

Asynchronous reuse of data is "rocket science" for many users, because they are accustomed to programs that use proprietary file formats (and operating systems with no real multi-tasking or **cron** capabilities). As soon as free standards are adopted, users begin asking for customizations and are willing to pay for anything that will increase their productivity. Moreover, free standards guarantee that customers are not making the wrong bet, as they won't ever be stuck with unusable data if the software market changes.

While the conventional model of software distribution concentrates all knowledge in a few companies (or one of them), open standards leverage technical knowledge to anyone willing to learn. Whereas a proprietary product can be supported only by a limited number of qualified consultants (whose number and quality is centrally managed), the number of consultants supporting a free software solution is virtually unlimited and the offer can quickly adapt to the request.


In a world where computers are just tools to accomplish some other goals, easy customization and quick maintenance are basic requirements of power users. In my opinion, free software will quickly gain the trust it needs to be a real market phenomenon. As soon as you start to trust some open-source products, you learn that they deserve more. GNU/Linux fans must be ready to offer support in order to fulfill the upcoming need for consultants.

Viability for Support Companies

Obviously, independent consultants don't cover all the needs of computer users. Several activities can't be handled by individuals. Red Hat and S.u.S.E. are demonstrating that creating and maintaining a distribution can be a good source of revenue even when the product is freely redistributable. Debian-based efforts are on the way, although less advanced—mainly because both Red Hat and S.u.S.E. bundled proprietary products with Linux in order to

The Best of Both Worlds

**InterBase 5
Now Available!**



InterBase 5 for Linux* The Perfect Development Combination


Now Linux developers have a REAL solution for their embedded database application needs – InterBase 5! Only InterBase 5 meets the needs of these applications with ease of installation, use and maintenance, robust functionality, a small footprint, and little or no administration. The result – extremely stable applications with very low total cost of ownership.

InterBase provides multi-platform, complete Java support, client and tools independence, and is remote access enabled.

With companies like Computing Options, Labor Ready, Motorola, Philadelphia Stock Exchange, Premier Data Systems, PrimeCare, and SAGE US enjoying high performance and scalability with InterBase, it's no wonder that the Linux community has been demanding it!

Call 800-451-7788 ext. 7182 today to order InterBase 5 for Linux, to request a **FREE CD-ROM** evaluation copy or for information on the InterBase Solutions Partner Program.

* Red Hat Linux Version 5.1



InterBase
SOFTWARE CORPORATION

© 1998 InterBase Software Corporation.

**Embed.
Deploy.
Relax.**

survive while the market share was low, while Debian is completely detached from proprietary products.

In addition to "creating and packaging" jobs, open-source companies can specialize in technical support, covering the situations where computer systems are of critical importance. Big business realities using computer systems in their productive environment won't be satisfied with either the external consultant or the in-house technician. They need to rely on an external structure that guarantees round-the-clock operation of their technological aids.

Even if GNU/Linux or any other operating system is demonstrated to be completely reliable, power users will need to rely on a support company as a form of insurance. The more important computers are for a production environment, the more people are willing to pay to be reassured that everything will go on working and to have someone "responsible" to call in case of any failure. Such a "power user" support contract could also include a provision for refunds in case of down time. Big support companies will be able to efficiently deal with it, and clients will be happy to pay high rates if they never need to call for assistance.

In short, I see no need for software companies to sell any product; the support environment is big enough to offer good business positions in Information Technologies. Those at the top could use some of the revenue to pay for free software development, thus gaining access to the best software before anyone else and associating their name with software products. As a matter of fact, this practice is already pursued by the big distributions.

Viability for Education Centers

Needless to say, schools and universities have the best interest in teaching information technologies using free software tools. Due to its technical superiority, free software environments have more to offer to the students, but also need more technical knowledge to be proficiently administered. I see no money saved here in choosing free operating systems over proprietary ones, but educational entities could better spend their money on hiring system administrators than on subsidizing some already-too-big commercial software company. While my country, Italy, is stuck with a few rules that offer more support for buying things rather than for increasing human resources, other countries are already moving in the right direction—Mexico and France, for example, have announced plans to use Linux in their public schools.

One more point leads toward free software in education: when students get jobs, they prefer to use tools they learned at school in order to minimize extra learning efforts. This fact should lead colleges to teach only those tools not owned by anyone—those that are libre. Schools should teach proprietary software only if two conditions apply: no viable alternative is available, and the company that distributes such software *pays* the school for teaching its product. Paying someone for a product and then freely advertising it for him is definitely nonsense.

Social Issues

A few social issues relate to choosing one software model over another one. Although I mark them as social, they have economic implications as well.

While free software is not cheaper than proprietary software if you bill for your own time, some environments use different rates in converting time to money. Most emerging countries have good intellectual resources but lit-

tle money, and they usually have many not-so-new computers as well. Proprietary operating systems are unaffordable for them, but free solutions are viable and productive. Actually, the "Halloween" document from Microsoft underlines that Linux is growing very fast in the Far East. Charity organizations usually have this same environment—little money and a good amount of human resources. This leads straight to the free software model for any IT requirement.

These ideas will probably suggest that free availability of information looks fairly leftist in spirit, as "information to the masses" looks quite similar to the old adage "power to the masses". What is usually ignored is the strong rightist flavour of the Open Source movement. The free software arena is fiercely meritocratic and a perfect environment for free competition, where the laws of the market ensure that only the best ideas and the best players survive. Proprietary standards, on the other hand, tend to diminish competition by decreasing innovation and consolidating previous results.

Limits of the Free Software Model

Naturally, I'm aware that not every software package can easily be turned into free software. I'm not talking about office products—I'm confident some good projects will supply this need, sooner or later. Rather, I'm talking about all environments where a strong competition exists for a product only loosely based on its software component. For example, industrial equipment might include a computer and some commodity hardware (a robot, custom I/O peripherals, PLCs, etc.); the software application hosted in the computer is a minor part of the whole, but its features greatly affect the overall value of the equipment. Producing and debugging such applications usually require huge investments (preventing free redistribution of the source code), as a form of protection against competitors.


Another meaningful example is cell telephones. They include a lot of software and such software is the component that defines the overall capabilities of the device. However, this software is almost invisible to the end user, who perceives the device as a telephone and not a computer. Such software is strictly proprietary because of its major functional role in the device.

Unfortunately, I see no easy way to liberalize this type of code. Although I don't care too much about cell phones (I don't use them), I would prefer to see free industrial applications because their technological content is usually worth reusing and adapting to new problems.■

Alessandro lives in one of the least Linux-aware towns in the least Linux-aware country in the world. He writes free software for a living and advocates free software for a mission. He hopes his upcoming child will keep off computers, recalling the good old times when such beasts were confined to their technical zoos. He reads e-mail as rubini@prosa.it, deleting spam and replying to everyone else.

Linux Projects Forum:
<http://www.linuxresources.com/apps/projects.html>

ONE IN A MILLION



All *Linux Journal* issues from 1997 on one CD-ROM
Order yours today, \$19.95 • 1-888-66-LINUX
<http://www.linuxjournal.com/>

P-Synch: Changing the Way We Change Passwords

Manufacturer: Mercury Information Technology, Inc.

E-mail: sales@m-tech.ab.ca

URL: <http://www.m-tech.ab.ca/>

Price: Customized—see web site

Reviewer: Tim Parker

One of the annoyances of administering networks is the need to change passwords regularly. We change passwords often for security reasons and that's fine. What bothers me is the need to log in to each machine individually and run the password changing program on each one, then log into each individual application that has its own user and password lists independent of the host machine and change the passwords there. My office network requires over thirty such changes; my home network, eight.

If I were running only UNIX, then I could use NIS and let that service change my machine passwords for me. However, NIS doesn't change non-UNIX passwords and doesn't do anything for application passwords. The same applies to Windows NT-based domains, where a central user list is maintained. Domain users don't extend to UNIX or applications. Utilities exist to provide some cross-platform support for NIS and NT domains, but I haven't found one that works well across my mixed network platforms.

Obviously, this is a problem that has plagued users for years. The folks at M-Tech (Calgary, Alberta) have done something about it. M-Tech's solution is called P-Synch (for password synchronization). It uses a number of scripts and API routines to perform all the password changes on your network from a single location. (It is an obvious solution, once you've seen it work.) I first encountered P-Synch two years ago in an earlier version and have used it religiously on my own networks, recommended it to many clients and written about it extensively since. A Linux implementation of P-Synch was an obvious spinoff of the UNIX version. Even better, M-Tech has released a new version of P-Synch which adds several new features that make it more functional and easier to manage.

As you may have gathered from the comments above, P-Synch uses scripts to change passwords on all applications and machines you tell it to access. P-Synch accomplishes this by using either a native API or TELNET to log in to each machine or application one at a time, running whatever commands are necessary to modify the passwords. You specify only the password to change to (or to reset) and P-Synch runs through its list of targets in the background. Since only a single copy of P-Synch is required anywhere on the network, no client programs need to be installed on each machine. The user interface can be character-based, GUI or HTML. An administrator defines which machines

and applications a user can modify passwords on, as well as more advanced options such as password aging.

Since P-Synch is script-based, it can change passwords on any machine or application that can be logged in to using TELNET or for which an API can be written. The list of M-Tech-provided scripts for operating systems is lengthy: Linux, many versions of UNIX, NetWare, Windows 95/98 and NT, LAN Manager, PathWorks, MVS and VMS. Application scripts available from M-Tech include Oracle, Sybase, SQL-Server, cc:Mail, MS-Mail, Exchange, GroupWise and GroupWare (including Lotus Notes). Scripts for other targets are starting to appear in Usenet newsgroups and on web pages, mostly written by P-Synch administrators.

P-Synch is fast, clean, easy to use and worth every penny.

The documentation for P-Synch may be downloaded from the M-Tech web site in PDF, PS or HTML formats. You can download uncompressed, ZIP or GZIP versions. The installation and configuration guide is about 260 pages long, but you will likely need to examine only a few pages to install P-Synch properly. Instead of printing the entire document, an Acrobat or PostScript viewer is best used to find those sections of interest. (Viewing on-line saves both time and paper.)

If you are running NIS (or a Windows-based alternative), P-Synch needs to run on the master. P-Synch will not function properly if installed on a client of an NIS master. (NIS does not allow administrative password changes from clients.) If you are not running NIS or a similar network-wide user management system, P-Synch can reside on any machine on the network. All machines must be running TCP/IP. P-Synch can coexist quite happily with NIS, handling the non-NIS targets only, if you prefer.

Prior to installing P-Synch, you must gather a list of the IP addresses of each machine on which P-Synch will change passwords. P-Synch requires root (or equivalent) access on each of these machines. It is handy to create a test account or two on the server and a few clients to make sure P-Synch is performing its tasks properly before trusting it to handle

your network-wide password management. These test accounts can be deleted after testing or saved for other purposes. A complete installation takes about 4MB of disk space (noticeably down from the last version's requirement of 9MB).

Installing P-Synch takes only a few minutes; however, configuring for a network with several applications and operating systems can take a while longer. The usual installation procedure is to copy or download the files to your Linux machine and extract the files in the library (tar, gzip or zip). One such library file is called `unix_srv.tar` and contains binaries for all supported UNIX and Linux platforms. After extracting this tar file, an installation shell script is run which handles the file setup procedure. A manual check of a configuration file to ensure it has the proper location for the `passwd` file (by default, it assumes `/etc/passwd`) completes the file setup. P-Synch normally uses TCP port 106, but this can be changed if port 106 is in use by another service. To test the installation, a TELNET session to the TCP port should produce a password prompt, at which point the installation is finished. The installation guide contains a list of common problems encountered when setting up P-Synch and they should account for most Linux system configurations.

P-Synch uses a script called `psynch.conf` to manage password changes. Separate parts of the program control users changing their own passwords, as well as root changing any password. On Linux and UNIX systems, P-Synch modifies passwords directly by interaction with the `passwd` program, not through a script (which would provide a potential security hole). The `psynch.conf` script can be edited if necessary, which makes it easy to handle special requirements such as firewalls and proxies as well as encryption schemes that manage password files. For non-UNIX passwords, P-Synch's `psynch.conf` file must be modified to use a script for password changes. M-Tech provides a number of prewritten scripts for different operating systems as well as applications that reside on UNIX or Windows machines (such as Oracle, Lotus Notes and so on). Non-UNIX systems require changes to the `inetd` file, but these can be cut-and-pasted from M-Tech's documentation or scripts. Linux and UNIX versions of P-Synch require a login called `psadmin`, which is used by the server to verify that P-Synch agents on other machines are allowed to change passwords. The `psadmin` login should be set up so that it has no access privileges.

Our test network consisted of three Linux machines (two Red Hat and one Slackware), four that were SCO OpenServer or UnixWare and three Windows NT servers, along with twelve Windows 95/98 machines. The server applications were Oracle 8, Lotus Notes, Exchange, Novell GroupWise and SQL-Server (all on the Windows NT servers). On this network, installation and configuration of P-Synch took about two hours. Most of that time was spent setting up the non-Linux/non-UNIX password change routines, with about half an hour required to debug the various scripts. If you are working on a Linux-only network, the process will take less than ten minutes.

Once properly configured, users anywhere on the network could run the P-Synch routines to modify their own passwords. The HTML interface in particular is friendly and attractive. Users can specify which machines or applications the change affects, or accept all (the usual case). Administrators can change passwords from any client. The amount of time required for a password change depends on the network load, the number of targets and the nature of

the operating systems. On our test network, the password changes went quickly, completing in under two minutes.

To test P-Synch with NIS, we set up one of the SCO machines as an NIS master and a Red Hat Linux system as the slave. We let NIS handle the password changes on half of the UNIX and Linux machines while P-Synch handled the rest of those types as well as the Windows machines. Propagation time for password changes didn't drop noticeably, which was expected since the Windows and application scripts are the major time consumers.

P-Synch is usually licensed to networks based on the number of users. M-Tech will customize the pricing plan to suit your requirements. Earlier versions of P-Synch cost about \$10 per user; the latest version is likely to be similarly priced. If you don't want to worry about the password-changing hassle anymore, you'll find P-Synch to be a wonderful utility. It is fast, clean, easy to use and worth every penny. The benefits are multiplied many-fold on heterogeneous networks. If you want more information about P-Synch, or want to download the application itself for evaluation, check out the M-Tech Web site at <http://www.m-tech.ab.ca/>. For the curious, a white paper describing P-Synch is available in Acrobat PDF and PostScript formats.■

Tim Parker lives in Ontario, Canada, and can be reached via e-mail at tparker@tpci.com. He is a widely published UNIX author with over 1,000 articles and 40 books on the subject. Dr. Parker's latest book is *Linux Unleashed, Third Edition* published by Sams. When not writing, Tim flies planes, scuba dives and argues with his temperamental network of thirty PCs and workstations. He often loses the arguments.



5th Annual
Linux Expo

**You've heard about the power...
Now feel the heat!**

Come join the crowd at the largest Linux event ever!
This is where the fire is! Linux Expo '99 will be the hottest ever, boasting a spicy blend of industry experts and exhibitors. Share the latest developments in Linux technology and it's blistering future.

**SIZZLING SPEAKERS
THRILLING TRADE SHOW
MULTIPLE CONFERENCE TRACKS
AND TUTORIALS**

**May 18-22, 1999
Raleigh, NC
Raleigh Convention Center**

Register NOW!
Phone: 919.383.0044 Fax: 919.383.0035
Email: linda@infomarketing-inc.com
www.linuxexpo.org

by David A. Bandel

Welcome to the second installment of my look at software packages worth downloading. I always find something good on the Internet. A lot of developers have been quite busy and are turning out some excellent software. The hot development library seems to be GTK+, and with so many novice Linux users who are not command-line oriented, it is a good thing.

gtkcookie:

<http://www.110.net/~pq1036/>

Do a lot of web surfing? Notice how many sites want to leave little droppings called "cookies" on your system? So many, in fact, I added a hard drive to store them all. Okay, I actually needed more space to download and compile programs like this one. It is a slick, easy way to view, edit and delete cookies. Frankly, I would rather eat them. Required libraries are gtk-1.0.6, Xi, Xext, X11, m and glibc.

guitar:

<http://disq.bir.net.tr/guitar/>

guitar is another of those utilities you love to have. I like this wrapper for tar despite the fact that I am quite comfortable on a command line. **guitar** lets you see what is inside a tar or gzipped tar file, read the text files inside, and extract the one you want. It also lets you create a directory for storing the files. This version does not create tar files, but I'll bet future versions will. Required libraries are gtk-1.0.6, Xi, Xext, X11, m and glibc.

gentoo:

<http://www.obsession.se/gentoo/>

gentoo is yet another file manager. It looks simple, but hides a lot of complexity. It will copy, move, and do other things as well. A three-row button bar on the bottom can be configured as you see fit. In fact, the most complicated part of this program seems to be figuring out how to fill up all the empty buttons rather than actually doing it. The configuration box takes a cue from other recent systems that use tabs across the top of the box to change to various configuration pages. Once it is configured (the default configuration is actually quite good), it is easy to use. Required libraries are gtk-1.0.6, Xext, X11, m and glibc.

slashes.pl:

<http://members.xoom.com/alexsh/slashes/>

slashes.pl is a small Perl script that all the "dotheads" in the audience will like. It is a way to keep the *Slashdot.org* headlines on your screen while having your web browser open to it all day long, just so you can press reload. This program allows you to see the headlines whenever you wish, just in case a news article appears that you would like

to read. While I had trouble with the "Read" button even after setting the line to point to my Netscape binary (it would crash), I suspect I just need to reread the script and find out what I am doing wrong for my setup. Required libraries are Perl, gtk-1.0.6 and the libgtk-perl module.

The Gaby Address Book of Yesterday:

<http://www.multimania.com/fpeters/gaby/>

Gaby is a deceptively simple address book. It carries all the latest fields for those needing to keep up with more than just phone and fax numbers. Space is reserved for URL, e-mail address and more. What it lacks is a way to connect to a database engine like MySQL. The flat-file method works well for individuals, but is not practical for a site where multiple users may have years of contacts listed. This one is great for single users and has the potential for larger sites if it can connect to a database engine. Required libraries are gtk-1.0.6, Xext, X11, m and glibc.

The Amazing Anagram Thingie:

<http://www.vis.colostat.edu/~scriven/anagrammer.php3/>

This is a command-line "game" to help you with those pesky anagram puzzles—very fast, very simple and easy to use. A phrase of any length will have this one scrolling anagram solutions off the screen for a long time. It is a great game for those puzzled by anagrams. Required library is glibc.

Ministry of Truth:

<http://tomato.nvgc.vt.edu/~hroberts/mot/>

This particular package has been out for a few months, has earned a place in my own system and will soon occupy an internal web site where I regularly work. Sometimes keeping track of jobs is difficult. This program makes it a breeze. It can track most users, software, equipment, jobs and people associated with them. While it requires MySQL, it handles everything once MySQL is installed and configured. For those requiring assistance, a (very low volume) user mail list has been set up. Required libraries are Apache with the php3 module compiled with MySQL, and MySQL.

Still lots more good packages out there. See you next month. ☺



David A. Bandel (dbandel@ix.netcom.com) is a Computer Network Consultant specializing in Linux. When he's not working, he can be found hacking his own system or enjoying the view of Seattle from an airplane.

LINUX TECHNICAL SUPPORT:

LINUX CONSULTANTS HOWTO: <http://consult.cyrius.com/lj.html>

The Future is in Your Hands

Back when you first started your mission, you kept it undercover. Your boss asked you to network the offices, and you made it look easy. Then you decided that the secret to your success shouldn't be a secret. So you took your OS public. And now that you are in control, there's no stopping you. There's no office that can't be networked, no server that can't be configured, no goal out of your reach.

NEW PRODUCTS!

Red Hat Linux 5.2

The latest version of the award-winning OS.

Red Hat Secure Web Server 2.0*

Control the security of web transactions, without the huge expense.

Linux Powertools

Over 200 extra packages of tools and software for Linux.

Check out these exciting new products @:

www.redhat.com

Red Hat® Linux.

Unlimited possibilities.

To place an order
call our sales office
at 1-888-REDHAT1

For all other inquiries
please call 1-800-454-5502.

Red Hat Linux v.



red hat®

5.2

© 1998 Red Hat Software, Inc. All rights reserved. "Red Hat" and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat Software, Inc. in the U.S. and other countries. The "Red Hat Secure Web Server" is a trademark of Red Hat Software, Inc. "Linux" is a registered trademark of Linus Torvalds.

*Due to restrictions imposed by the US Government, this version of The Red Hat Secure Web Server may not be shipped outside the United States and Canada, or to a national of a country other than the United States and Canada. Thank you for your cooperation.



Caching the Web, Part 2

This month Mr. Guerrero tell us about the definitive proxy-cache server, Squid.

by David Guerrero

Last month we discussed the basic concepts of proxy servers and caching. Now, let's see how to implement this technology in your organization. A few proxy-server programs are on the market, such as MS-PROXY, aka Catapult, available only for Windows NT, and Netscape Proxy Server, available for different UNIX platforms and Windows NT. Both have two main drawbacks: they are commercial software and they don't support ICP. The excellent Apache web server has included a proxy-cache module since its 1.2 version. This module is a very interesting option: it's free, and works with the most popular web server on the Net. However, it doesn't use ICP, and its robustness is not comparable to the best choice for a proxy-cache server—**Squid**.

Squid is a high-performance proxy-cache server derived from the cache module of the Harvest Research Project, maintained by Duane Wessels. It supports FTP, gopher, WAIS and HTTP objects. It stores hot objects in RAM and maintains a robust database of objects in disk directories. Squid also supports the SSL protocol for proxying secure connections and has a complex access control mechanism. Another interesting feature of Squid is negative caching, which saves "connection refused" and "404 Not Found" replies for a short period of time (usually five minutes).

Squid consists of four programs:

- **squid**: the main proxy server
- **dnsserver**: a DNS lookup program that performs single, blocking DNS operations
- **unlinkd**: a program to delete files in the background from the cache directory

It also provides a CGI program, designed to be run through a web interface, that outputs statistics about its configuration and performance and allows some management capabilities.

Squid Installation

Installing Squid is easy. Just download the source archive from <http://squid.nlanr.net/> and, in a temporal directory, type:

```
gzip -dc squid-x.y.z-src.tar.gz | tar xvf -
```

Next, compile and install the software by typing:

```
cd squid-x.y.z
./configure
make all
make install
```

These commands install all needed programs and configuration files to `/usr/local/squid`. The binary programs are installed in the `/bin` directory, the configuration files in

`/conf`. Log files are located in the `/logs` directory, and the object database in the cache directory and its subdirectories. A shell script called **RunCache** is in the `bin` directory used to run the squid binary, and assures that if the process dies for any reason, it is restarted automatically. So, put the following line in your `rc.local` file:

```
/usr/local/squid/bin/RunCache &
```

Squid is a high-performance proxy-cache server that supports FTP, gopher, WAIS and HTTP objects.

This will generate an error log in `/usr/local/squid/squid.out`, if Squid could not start because of some configuration problem.

Of course you can choose to install an RPM version of Squid if you use RedHat Linux or another distribution that supports RPM packages.

Squid installs a sample configuration file called `squid.conf` with many comments for each option. Here you can change the ICP and HTTP ports (3128 by default) and define how much memory and disk space to reserve for caching objects and other parameters such as refresh patterns and access control restrictions. Of course, you need an ICP port only if your cache is going to be the sibling or parent of other caches. The directives for changing these values are `http_port`, `icp_port`, `cache_dir` and `cache_swap`. Additionally, you can set the maximum object size to be stored in the database; the default is 4MB. Also, you should uncomment the following lines in this file:

```
cache_effective_user nobody
cache_effective_group nobody
```

This avoids running Squid as root, a dangerous habit for anyone who runs servers like **httpd** or **gopherd**. If you are using a recent version of Squid (at the time of this writing, the current version is 1.1.16), it will not start running as root, but will write an error message to the `squid.out` file.

To let Squid use 100 MB of your HD, the directive `cache_dir` should be something like this:

```
cache_dir /usr/local/squid/cache 100 16 256
```

Before starting Squid for the first time, create the cache and logs directories. To build the cache and hashed subdirectories, you should execute the commands:


```
cd /usr/local/squid
mkdir cache
chown -R nobody cache
cd /usr/local/squid/bin
./squid -z
```

Finally, to create and change the owner of the logs directory:

```
cd /usr/local/squid
mkdir logs
chown nobody logs
```

Now Squid can be run safely for the first time, with the above RunCache invocation. It will spawn several **dnsserver** processes and write its PID in the file `logs/squid.pid`. Important warning or error messages can be found in the `squid.out` and `logs/cache.log` files. Remember, if you want to shut down the cache, you must first kill the RunCache process to avoid an immediate restart and then type:

```
/usr/local/squid/bin/squid -k shutdown
```

Never use `kill -9` to shut down the cache, because it doesn't close the object database in such a way that it can be recovered—you'll probably lose part of it.

Restricting Access to Your Cache

In order to enable only those users who are in your organization to access your cache, you must set up some access control lists (ACLs). Defining access lists in Squid is quite easy; all access lists are defined with a name and are used to define a subset of elements. You can make a subset of IP addresses, protocols, destination URLs and even browser brands. The directive to define an ACL or subset is:

```
acl name type data
```

You can learn more about ACL types in the example `squid.conf`. In the case of restricting access to only our users, the type needed is `src`. For example, suppose you want to allow access to the cache to all browsers in the 172.16.236.0 class C, the first 32 addresses of the next class C and your PC, 172.16.237.180. You can define an ACL like this:

```
acl my_users src 172.16.236.0/255.255.255.0
acl my_users src\
172.16.237.1-172.16.237.32/255.255.255.255
acl my_users src\
172.16.237.180/255.255.255.255
```

Next, define an ACL for the rest of the addresses. This line is included in the `squid.conf` example file:

```
acl all src 0.0.0.0/0.0.0.0
```

Apply these ACLs in an ordered way with the `http_access` directive. The syntax is:

```
http_access allow/deny [!]acl .. [!]acl
```

For example:

```
http_access allow my_users
http_access deny all
```

More than one ACL can be combined in the same `http_access` directive and can be used in its negative form (i.e., preceded by `!`). The example shown is the most simple use of ACLs, but more complex forms will allow connections only in designated hours and days, allow only defined URLs or domains to be fetched and restrict some protocols such as FTP. This powerful feature of Squid can help you enforce and implement your security policy, whether you use Squid

SW TECHNOLOGY

Up to 600MHz **Linux/Alpha** boxes
Single and Multi-processor
Linux/Sparc stations

Dual and Quad Pentium Pro Processor Servers Running SMP Linux

Up to 1000MB RAM and 512k on-chip L2 cache

Linux Laptop computers with 800x600 or higher
X resolution on the LCD
Linux PCMCIA support
Linux and Windows 95 dual boot and worldwide
warranty services

NCR 53C810 PCI fast SCSI II controller for \$53.00
(No, that is not a misprint)
Other NCR controllers also available

Pentium/Pentium Pro Linux Machines
customized to any flavor you like
Best available components and prices

You will find all of the above by pointing
your web browser to

<http://www.swt.com>

where you can also find how Java does
real work—to provide you dynamic pric-
ing and configuration information as you
browse

SW TECHNOLOGY

1701 N. Greenville Avenue Suite 1104
Richardson, TX 75081, USA
phone: 972-907-0871
fax: 972-907-9339
info@swt.com

in your firewall or the Squid machine is the only one allowed to cross your firewall. Just look for examples in `squid.conf`.

There is also an ACL to permit setting the desired web ports you allow your users to use. This is the `Safe_ports` ACL. You should uncomment this line and add the 443 port to this ACL in order to allow the use of secure web servers through your Squid server.

A Look at the Logs

Squid can generate huge logs of your proxy-cache usage. With this information and the help of some scripts, we can generate complete access statistics, like the ones generated from web servers. Squid maintains three main log files:

- **cache_log** includes warnings and information about the status and operational issues of the cache.
- **store_log** includes information about database operations, such as inserts of new items and releases of expired objects.
- **access_log** contains an entry for each object fetched from the cache and information on how it was served. It also includes information about each ICP query received by the cache from other servers using this server as a neighbor.

Many utilities are available for generating statistics from the `access_log` file (see Resources).

Remember, it is not considered ethical to surf your `access_log` to see which places your users visit. Some sites have chosen not to publish processed statistics in any form to guard their users' privacy, which is an important concern for all of us involved in the Internet community.

The logs grow very quickly and in a few days can eat up your remaining disk space. To safely clean your log files, you should rotate them with the `SIGUSR1` signal. A single line can be added to your `crontab` to begin new log files each night:

```
/usr/local/squid/bin/squid -k rotate
```

This command will create the files `access_log.0`, `store_log.0` and `cache_log.0` and begin logging to new empty files. Now you can safely remove these files or process them for statistical purposes. The next time you rotate

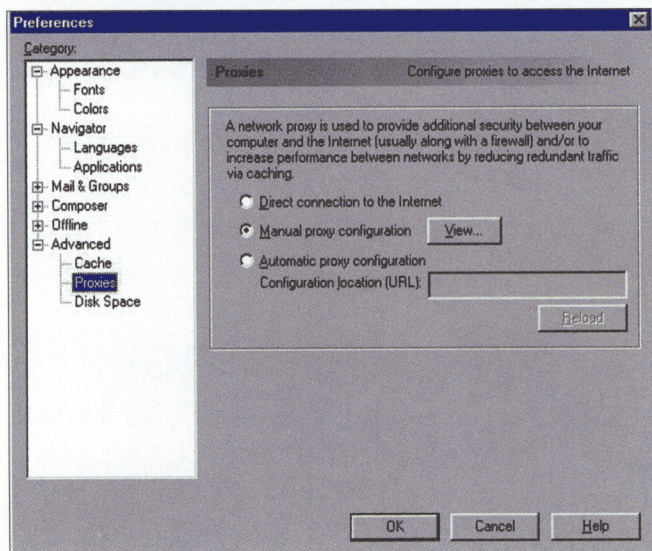


Figure 1. Proxy Preferences Screen

the logs, `files.0` will be moved to `files.1` and so on. You can configure how many extensions Squid will use for these rotations to save disk space with the `logfile_rotate n` directive in the `squid.conf` file.

Configuring Browsers to Use Cache

To begin using your new proxy-cache server, you must first instruct your user's browsers to fetch objects from your server instead of retrieving them directly. In most modern web browsers, one of the configuration options is the specification of the proxy setup. Another option is to specify a list of domains or URL patterns which must be fetched through the proxy.

In Netscape Navigator or Communicator, you can include a proxy server and its port for each service to be proxied. With Squid, you can use these settings for the HTTP, Security (SSL), FTP and WAIS services, all with the same port (3128, by default). First, select the "Manual proxy configuration" radio button and then the "View" button to type in your settings. Figures 1 and 2 show examples of these screens.

Another solution is the Automatic Proxy Configuration, introduced in Netscape

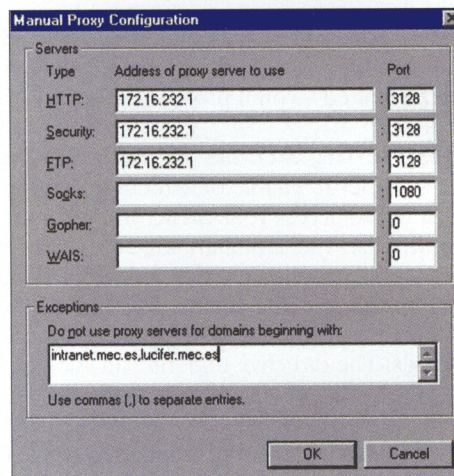


Figure 2. Manual Proxy Configuration Screen

Navigator 3.0, that allows multiple proxy servers, backup servers and different servers by domains. This configuration sits in a Javascript-like file that must be retrieved from a server. Using it, you can change the topology of your cache mesh or introduce new servers that must be treated as "No proxy for" servers. Without telling your users to change their configurations, the new configuration script is reloaded each time the browser is launched. MS Internet Explorer has also supported the automatic proxy configuration feature since version 3.02.

An example of this kind of configuration for Netscape Navigator and Communicator is shown in Figure 3. In this example, each time the browser is started, it loads the file `proxy.pac` from the server `intranet.mec.es`. This file must be returned with MIME-type `application/x-ns-proxy-autoconfig` which can be accomplished in two ways:

1. Add the following line to your `mime.types` file:

```
application/x-ns-proxy-autoconfig      pac
```

2. Or add the following line to your `Apache srm.conf` file:

```
AddType application/x-ns-proxy-autoconfig pac
```

For the changes to take effect, you must name your proxy auto-configuration file with the `.pac` extension and restart your web server. The Netscape documentation will tell you about the syntax of the `.pac` file (see Resources).

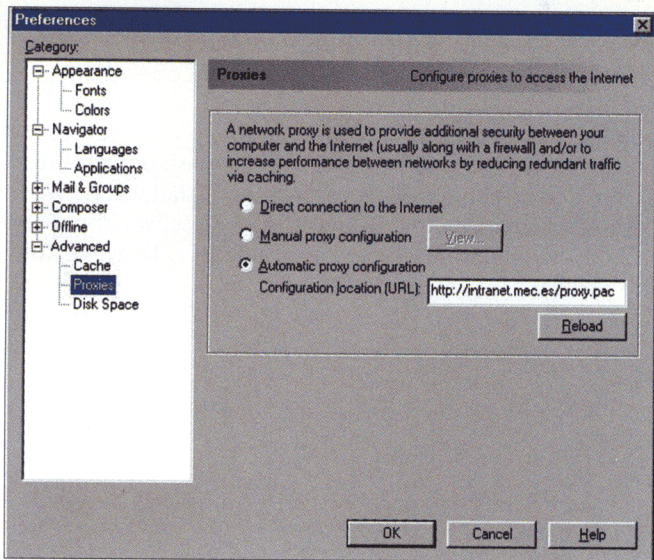


Figure 3. Automatic Proxy Configuration Screen

Nevertheless, we'll look at a couple basic examples of how to write them.

No HTML tags should be embedded in the Javascript file, just the function **FindProxyForURL** with arguments *URL* and *host*. This function should return a single string containing **DIRECT** (get the object directly from the source), or **PROXY host:port** (get the object through this server and port). The string can contain more than one of these directives, separated by semicolons. For example:

```
function FindProxyForURL(url, host)
{
    return "PROXY proxy1.mec.es:3128;
    PROXY proxy2.mec.es:80; DIRECT ";
}
```

will instruct the browser to use the first proxy to fetch the object. If it can't contact the first (proxy1), then it will try the second (proxy2); in the case that both are down, it will fetch the object from the source. This gives a fault tolerance level to our cache system.


One interesting feature is using different proxies for different domains and including support for internal servers where we don't want to use the cache. For example:

```
function FindProxyForURL(url, host)
{
    if ( isPlainHostName(host) || dnsDomainIs(host,
        "intranet.mec.es"))
        return "DIRECT";
    else if (shExpMatch(host, "*.com"))
        return "PROXY proxy1.mec.es:3128";
    else
        return "PROXY proxy2.mec.es:80";
}
```

This function will directly fetch all objects whose URL is only a word with no dots or the Intranet server, all .COM objects from proxy1 and the rest from proxy2.

As a tip, the .pac file can be generated "on the fly" by a CGI script, giving different proxy configurations for different browsers, e.g., depending on the REMOTE_HOST environment variable provided by the CGI interface. In this way, load balancing between different networks can be achieved.

THREE REASONS TO USE COMPUTONE FOR REMOTE ACCESS & CONNECTIVITY SOLUTIONS RAW POWER RAW POWER & MORE RAW POWER



IntelliServer
921.6 KBPS



Linux Friendly



IntelliPort II EX
Intelligent Processor

COMPUTONE
Corporation

FOR MORE INFORMATION CONTACT:

1-800-241-3946 x 1616 or 770-625-0000 x 1616

linux@computone.com ▲ www.computone.com

Resources

Squid Home Page: <http://squid.nlanr.net>

Squid FAQ: <http://squid.nlanr.net/Squid/FAQ/FAQ.html>

ICP Draft: <http://ircache.nlanr.net/Cache/ICP/ICP-id.txt>

Squid users mailing list: squid-users-request@nlanr.net with "subscribe squid-users" in the body.

Squid announce mailing list: squid-announce-request@nlanr.net with "subscribe squid-announce" in the body.

Automatic Proxy Configuration File Format: <http://home.netscape.com/eng/mozilla/2.0/relnotes/demo/proxy-live.html>

Squid Log File Analysis Tools

Cache Stats, by Iain Lea:

<ftp://ftp.ecrc.de/pub/www/cache/squid/contrib/cache-stats/>

Squid Clients, by Nico Tranquilli:

<http://www.cineca.it/~nico/squidclients.html>

Squid Times, by Nico Tranquilli:

<http://www.cineca.it/~nico/squidtimes.html>

Calamaris, by Cord Beerman:

<http://www.hoexter.netsurf.de/homepages/cord/tools/squid/>

Always remember that the MIME-type returned by the CGI must be `application/x-ns-proxy-autoconfig`.

Joining a Hierarchy

If your cache is to be part of a cache mesh or your proxy server is to be connected to another proxy that will be its parent, you must use the `cache_host` directive. You must include one line for each of your neighbors. The syntax for this line is:

```
cache_peer hostname type http_port icp_port options
```

where:

- *hostname* is the name of your neighbor.
- *type* is one of parent or sibling.
- *http_port* is the neighbor's port from which to fetch objects.
- *icp_port* is the port to which ICP queries are sent. Use a value of 0 if your neighbor does not run ICP, or 7 if your neighbor runs the UDP echo service. This can help Squid to detect if the host is alive.

Squid includes a simple, web-based interface to monitor the cache performance and provide useful statistics.

You can specify the option `default` to use this host as a last resort in case you can't speak ICP with your parent cache. Another option is the `weight=N` to favor a specific parent or sibling in the neighbor selection algorithm. Larger values give higher weights.

If you have a stand-alone cache, you should not include any of these directives. If you have one parent that runs its HTTP port on 3128 and its ICP port on 3130, the line to include in the `squid.conf` file is:

```
cache_peer your.parent.cache parent 3128 3130
```

With the `cache_peer_domain` directive, you can limit which neighbors are queried for specific domains. For example:

```
cache_peer_domain first.parent.cache .com .edu
cache_peer_domain second.parent.cache .es .it .uk .fr
```

will query the first cache only for the `.COM` and `.EDU` domains, and the second for some of the European domains.

If you have only one parent cache, the overhead of the ICP protocol is unnecessary. Since you are going to fetch all objects (HITS and MISSES) from the parent, you can use the `no_query` option in the `cache_peer` directive to send HTTP queries to only that cache.

Also, there are some domains you will always want to fetch directly rather than from your neighbors. Your own domain is a good example. Fetching objects belonging to your local web servers from a faraway cache is not efficient. In this case, use the `always_direct acl` command. For example, in our organization we use:



▲ Leverage Information Systems, Inc.

This train flies!

The Locomotive Web Application Server from Leverage.

Since 1996 we've been building the Locomotive to deliver unsurpassed speed, power and scalability for the consumer market.

Version 4.0 is
tuned to fly
on Linux.



Next stop, Open Source!

Find it at <http://www.locomotive.org>

Leverage Information Systems
www.leverage.com

220 Sansome St., Suite 1090

San Francisco, CA 10904

ph: 415 512 9700 fax: 415 512 9701

e-mail: info@leverage.com


```
acl intranet dstdomain mec.es
always_direct allow intranet
```

to avoid getting our own objects from the national cache server.

The Cache Manager

Squid includes a simple, web-based interface called **cachemgr.cgi** to monitor the cache performance and provide useful statistics, such as:

- The amount of memory being used and how it is distributed
- The number of file descriptors
- The contents of the distinct caches it maintains (objects, DNS lookups, etc.)
- Traffic statistics with each client and neighbors
- The "Utilization" page, where you can check the percentage of hits your cache is registering (and thus bandwidth you are saving).

Be sure to copy the **cachemgr.cgi** program installed in your `/usr/local/squid/bin` (or wherever you chose) to your standard CGI directory, and point your browser to `http://your.cache.host/cgi-bin/cachemgr.cgi`. There, you should type your cache host name, usually "localhost" or the name of your system, and the port your cache is running, usually 3128, and check all the options.

Conclusions and Tips

A proxy-cache server is a necessary service for almost any organization connected to the Internet. In this article, we have tried to show the whys and hows to implement this technology, and a brief tutorial on Squid, the most advanced and powerful tool for this purpose. Don't forget to read all the comments in the example configuration file. They are complete and useful and show a lot of features not mentioned in this article.

Perhaps in a few years, with the growth of PUSH technology and the use of dynamic content on the Web, caching won't be a solution to the bandwidth crisis. Today, it's the best we have.

One problem proxy caches don't solve is making certain your users configure their browsers to use the caches. Users can always choose to bypass your proxy server by not configuring their browsers. Some organizations have chosen to block port 80 in their routers except for the system

running the proxy-cache server. It's a radical solution, but very effective.

Another thing you can do to improve the speed of your users' browsers is pre-fetching the most accessed web sites from your cache. Recursive web-fetching tools which support proxy connections can help do this task in non-peak hours, e.g., **url_get**, **webcopy**. Launching one of these retrieval tools with the standard output redirected to `/dev/null` updates the cache with fresh objects.■




David Guerrero is a system and network manager for the Boletín Oficial del Estado. He has been using Linux

since the .98pINN days and now is playing with some Alpha-Linux boxes. When not working or studying, he likes to spend time with his love Yolanda, travel, play guitar and synths, or go out with his "colegas". He can be reached at david@boe.es.

www.linuxjournal.com



You Will Be
performance
reliability
service
Satisfied



everything you expect... and a little more



<http://www.linux-hw.com>
1-888-linux-hw

Visit us in the Linux Pavilion at Comdex 1998

Linux Hardware Solutions and the Linux Hardware Solutions Logo are trademarks of Linux Hardware Solutions, Inc. Red Hat and the Red Hat Logo are registered trademarks of Red Hat Software, Inc. "Linux" is a registered trademark of Linux Torvalds.

KDE: The Highway Ahead

Mr. Dalheimer describes some of the plans being made for future versions of KDE.

by Kalle Dalheimer

The K Desktop Environment (KDE, see <http://www.kde.org/>) has already generated a lot of interest, and many individual users and institutions alike are using it as their desktop environment of choice. However, nothing is so good that it doesn't have room for improvement. Since the beginning, we have considered stability more important than sticking to announced dates, so there have been occasional delays when we considered KDE not stable enough for a release. That said, we hope to release KDE 1.1 after one or two beta versions by the end of 1998. This version will contain a number of bug fixes, the much sought-after support for ICQJava, more robustness in the HTML display code and a few new features such as configurable key bindings for the whole desktop. Originally, we had planned to release only bug fixes as version 1.0.1. However, the completed new features have been requested many times and we don't expect to have 2.0 stable for some time, so we want to include everything that is ready. This way, we can move on at full steam after 1.1 is released.

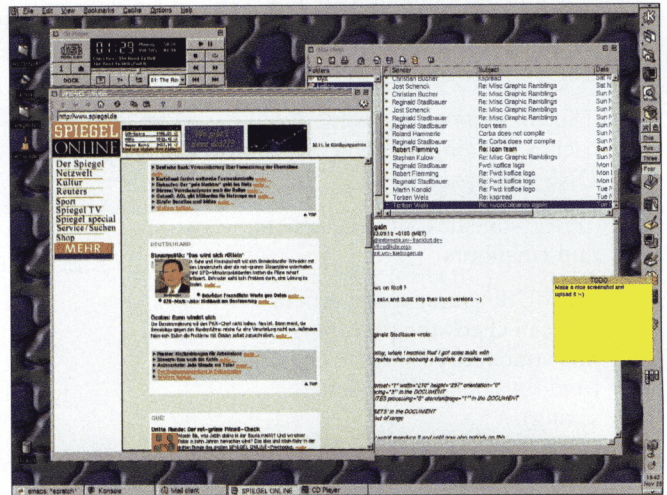
KDE 2.0, the next "official" version, will probably be out in late summer 1999, but this is still uncertain. In the meantime, we will also release a first alpha version of KOffice (discussed below).

Planned Features for KDE 2.0


We plan substantial rewrites and reorganizations of KDE for version 2.0. This will probably lead to snapshots that are unstable or don't compile, but as always, KDE follows the open-source model closely and makes daily snapshots available via the FTP site and the CVSup server. (CVSup is a software package for distributing and updating source trees from a master repository on a remote server host.)

Some of the new code that will go into KDE 2.0 belongs to the area file manager/web browser/networking code. Most of the new code is already written and just waiting to be committed to the CVS after the release of KDE 1.1. Changes include a complete rewrite of the file manager and the networking code which provides a great improvement over the current version, the ability to browse compressed archives and a complete overhaul of the user interface of the file manager. Related to this is the HTML widget currently in the process of being reorganized to make it more adaptable to new HTML standards and to provide better support for JavaScript and a free Java virtual machine. (**kaffe**, <http://www.kaffe.org/>, is a candidate if it supports standard AWT, Abstract Window Toolkit, programs by then.) Most of this is not just vaporware or simply plans for the future, but existing code, part of which is also available via the CVSup server with anonymous access.

Another program to be completely rewritten is the mail program KMail. The new version, dubbed KMail 2, will be more robust and flexible with respect to various attachments and feature IMAP support.



KDE 1.1 Desktop



KDE 1.0 Desktop

Of course, there is not only revolution such as complete rewrites, but also evolution as seen in the continuous development of existing features. More and more configuration modules are being added to the control center and we plan to provide more modules geared specifically toward system administrators. We hope to be able to support both LinuxConf and COAS (Caldera Open Administration System), but are definitely in need of more skilled developers to help in this area.

Another area of improvement is the drag-and-drop protocol where we will switch from our "homegrown" protocol to the XDND protocol, since it will likely become the future drag-and-drop standard on (at least free) UNIX systems and is currently supported by programs written with the JX toolkit. The GNOME developers have indicated that they will also support this protocol, possibly leading to



Your One Stop Internet Hardware Source

"Telenet System Solutions web server produced the most surprises during our tests, with a BSDi-powered, single-CPU system that kept up with-and in some cases outperformed- twin-CPU machines running Windows NT from HP, Micron and Dell."

- By Sean Fulton, May 05, 1997, Issue: 661 **Communication Week.**

Telenet System Solution's offer state of the art Internet Servers, RAID Servers, PC work-stations, WAN Products and peripherals at competitive price. At "tesys.com" we'll help you find the right system hardware, configure it, price it, build it and test it with OS of your choice. That's what you get when the company that sells you the complete system is also the one which custom builds it for you. Please visit our website or call us with your specific requirements.

New Products

TELEPRO-200/250 (2U/2.5U)




TELEPRO-500 Series (5U)



CROWN-440 RAID Server (4U)



INTERNET SERVERS



ISP Turnkey, LAN/WAN Gateways
WEB, FTP, DNS Servers
Proxy/Firewall, News Servers
Sun/Sparc (Ultra Axi) Servers
Intel/Alpha Servers

**Linux, BSDI, FreeBSD,
Microsoft NT, Solaris**


ENG. WORK STATIONS



Pent, P-II and Xeon based, high speed/
graphics, Multi-media, CAD / CAM
Systems for developers and graphics
designers. Support for Microsoft
Win95/NT, X-Windows
(Xi Graphics, XFree86, MetroX)

**Multi-boot operating systems &
SMP support available.**

RAID PRODUCTS



RAID 0, 1, (0+1) & 5,
(Rackmount/Towers)
Software & Hardware based RAID
SCSI to SCSI, PCI to SCSI Solutions
Removable hotswap RAID cards
Removable/Hotswap HDD, PS, Fans

**DPT, Adaptec, CMD, Mylex,
Kingston**


SLIM SERVERS



Specially Designed Space Saver Servers
ATX and NLX Form Factors
Rackmount/Desktop Models available
2U, 2.5U, 3U Configurations

**Featuring 3.5" high (2U) full
featured Rackmount Server,
Space Saver Series**


LAN/WAN PRODUCTS



CSU/DSU
Terminal Server
Multiport Serial Cards
ISDN/T1-Routers
Rack Modem Pool
Hubs / Switches

**Livingston, Ascend, Cisco, BAT,
Adtran, SDL, Compaq, Cyclades,
Digiboard, 3Com/USR**

RACKMOUNT PRODUCTS



Racks, Drawers, Slides, Keyboards,
Monitors, UPS's, Single board Computers.
AC/DC options. Hot swap PS options.
Single Keyboard, Monitor, Mouse-Digital
Switch Boxes.

**A Large variety of enclosures,
various sizes, AT/ATX.**

<http://www.tesys.com>

E-mail: sales@tesys.com

2480 Kruse Drive, San Jose, CA 95131 • Phone (408) 383-0334 ext. 130 • Fax (408) 383-0335

Toll-Free 1-877-7-Telenet (783-5363)

a first point of interoperability between the KDE and GNOME programs.

We will also make use of some new features provided by version 2.0 of Qt, the toolkit used with KDE. Among these are themes and Unicode support. Themes are something I personally consider less important for a desktop aimed at improving productivity, but I know a lot of users want it and we aim to please. Note that some theme ability is currently in KDE (see e.g., <http://kde.themes.org/>), but the new code will extend this to single widgets within the application windows.

Plans have been made to tear the window manager engine out of the window manager KWM and put it into the library, so that there can be different window managers to implement a different look-and-feel and still provide the window manager functionality needed by a full-featured KDE desktop.

This version will contain the much sought-after support for ICQJava, more robustness in the HTML display code and a few new features.

Unicode (see <http://www.unicode.org/>) is very important to gain acceptance in the Far East and other areas of the world with scripts containing more than 256 characters. Two Chinese localizations of KDE already exist, but these require a patched X server that combines two character codes transmitted for display into one. With Unicode support, such tricks will not be necessary, as the KDE message files can then contain 16-bit characters directly.

The usability of Unicode support is based heavily on availability of decent fonts for the script in question, something UNIX systems have traditionally been lacking. I have been looking at integrating the free True Type engine Freetype into KDE, but this is still in the beginning stages. (Contributions are welcome.) Another option is using a font server that supports True Type fonts.

Another area where work is being done is making KDE accessible for handicapped users. It is already possible to use very large fonts with KDE programs and especially to set such a large font once and for all for all KDE applications, so that people with slight vision impairment can use KDE programs, but this is not enough. We have had some success with voice-type software, i.e., software that allows users to navigate and operate the desktop and applications by speaking commands into a microphone. We are working with universities on leveraging their research results and making usable programs out of them. Another feature that falls into this category but has not been addressed yet is screen readers, i.e., software that reads out whatever is visible on your screen via your sound card and the speakers. While the necessary text-to-speech synthesis is a non-trivial problem (even though programs such as **emacspeak** show that it can be done in free software), screen readers become even more difficult to write when graphical user interfaces are involved, because a single text flow is no longer available as on terminal screens.




Figure 1

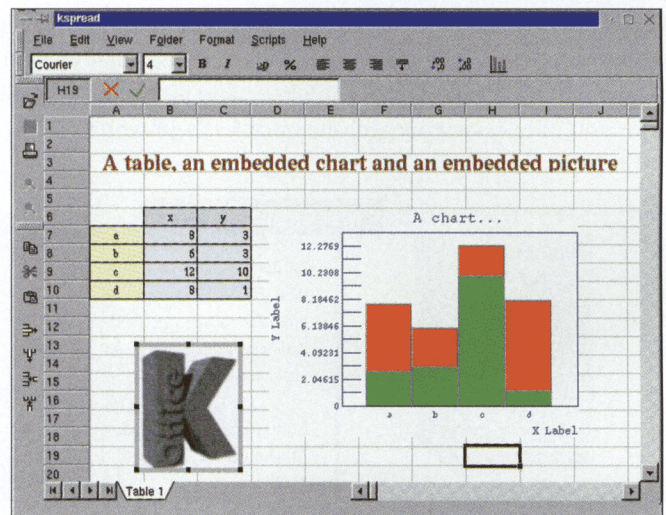


Figure 2

KOffice

The topic most interesting to readers is the development of KOffice. KOffice (see <http://koffice.kde.org/>) is intended to be a complete office productivity suite like Microsoft Office or Lotus SmartSuite. It is CORBA-based (using the high-quality, freely available ORB Mico that makes good use of today's compiler technology) and uses the Open Parts as its object model. Open Parts is the master's thesis topic of one of our developers and provides a rich object model on top of the rather bare-bones CORBA standard, including addition of event services and event filtering to CORBA. Open Parts can even be used with non-KDE applications.

Not only the large KOffice applications will use CORBA, but this technology will also be used in other areas, including the panel which would then be better able to swallow other programs. On the other hand, adding CORBA compliance to a relatively small program like the panel means bloating it, and we still have to investigate whether this is truly worth it.



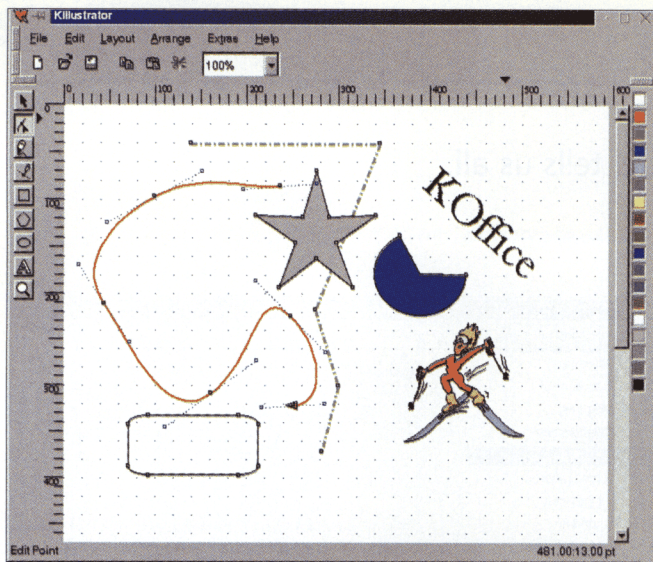


Figure 3

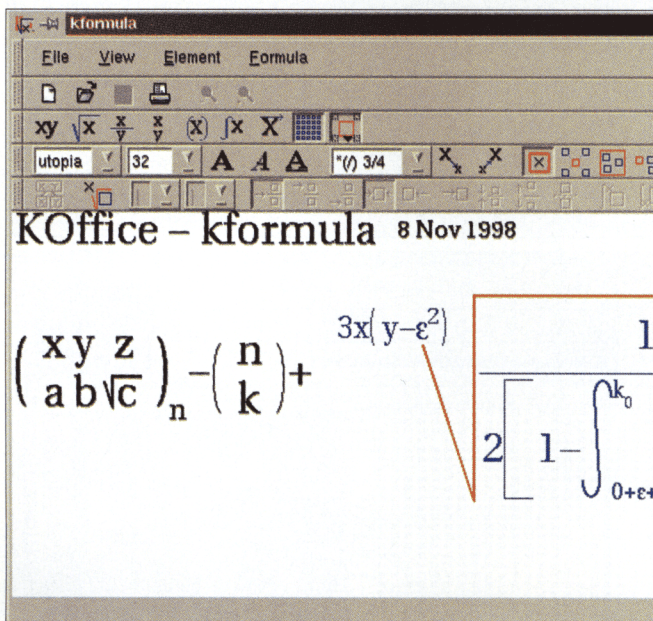


Figure 4

Currently, KOffice consists of the following applications which, as I write, have different levels of usefulness. First, there is KPresenter (see Figure 1), a presentation program in the spirit of MS PowerPoint that has already been used for "real work" (like my KDE presentation at the last International Linux Congress in Cologne) and features a large number of blending effects. Then, there is KSpread (see Figure 2), a spreadsheet that already contains about 20% of the functionality of MS Excel and is easily extensible via scripts written in Python.

A recent addition to KOffice is KWord, a word processor aimed at smaller size documents like personal letters, as well as complete books. Besides normal word processing features, it also has some features known from DTP programs, such as Adobe FrameMaker. For large documents or those containing many formulas, it might be better still to use KLyX, a WYSIWYG (what-you-see-is-almost-what-you-get) document processor that uses LaTeX as its formatting engine and is based on its older brother, LyX. KLyX is not yet as fully integrated into the KOffice concept as the other applications, but hopefully it will be in the near future.

The KOffice application that is probably most advanced at the moment is KIllustrator (see Figure 3), a vector-based drawing package. It features a lot of drawing tools and is very usable. Two other applications of KOffice are KFormula (see Figure 4), a mathematical formula typesetter; and KImage, a smaller image-manipulation program. Finally, there is KChart, a component for drawing business graphs such as bar charts and pie charts; and KOrganizer, an organizer program in the spirit of Lotus Organizer that is quite usable and sports many features and different views.

Since all KOffice programs are based on the Open Parts framework, they can all be embedded into each other. You can nicely embed a spreadsheet document into your word processor document and beef the whole thing up with a vector drawing done with KIllustrator. This embedding also features in-place activation with dynamic menu bars, and unlike other embedding technologies on some platforms we like to hate, it is quite stable.

We are very aware that for Linux to become a viable alternative to Windows or the Macintosh, it must feature a consistent easy-to-use desktop and also have an office productivity suite that need not (or should not) be as bloated as Microsoft Office, but definitely needs the features used on a daily basis and must be able to read common document formats. Unfortunately, Microsoft Office formats are not well-documented. Still, with the help of the work done in LAOLA (library for Microsoft structured storage files), we hope to be able to support at least some of them. Even though I think this is a boring task (if somebody thinks I am wrong here, you are welcome to help with this project), we will also write an RTF reader and writer, and I have already successfully imported some FrameMaker documents via my homegrown filter into KWord. Whether other formats (such as WordPerfect, StarOffice or the Lotus SmartSuite) will be supported depends on two things: whether the respective manufacturers will make their format documentation available and whether someone volunteers to do the work.

Summary

As you can see, many things are happening in the KDE world and development is occurring faster than ever. As with all free software projects, we can be only as good as the people who support us, so if you are interested in GUI hacking, writing file-format filters or any of the other topics I have mentioned (or something completely different that fits into the overall KDE concept), don't hesitate to contact us (see <http://www.kde.org/> for contact information). This is truly an exciting project to be in.

Kalle Dalheimer is a freelance software consultant and a technical writer/technical editor for O'Reilly. He is a member of the KDE core team, where he is in charge of the libraries and some of the applications. When not hacking or writing, he plays with his two-year-old son, setting up wooden miniature railway systems. He can be reached via e-mail at kalle@dalheimer.de.

COOL
it works with
LINUX

**LET THE WORLD KNOW YOUR PRODUCT
 WORKS WITH LINUX:**
<http://www.linuxjournal.com/cool.html>

The GNOME Project

What is GNOME and where is it heading? Miguel tells us all.

by Miguel de Icaza

GNOME is an acronym for GNU's Network Object Model Environment. GNOME addresses a number of issues that have not previously been addressed in the UNIX world, such as:

- Providing a consistent user interface.
- Providing user-friendly tools and making them powerful by leveraging the UNIX foundation.
- Creating a UNIX standard for component programming and component reuse.
- Providing a consistent mechanism for printing.

GNOME's main objective is to provide a user-friendly suite of applications and an easy-to-use desktop. As with most GNU programs, GNOME has been designed to run on almost all strains of UNIX-like operating systems.

History of GNOME

The GNU GNOME project was initially announced in August 1997. After just one year of development, approximately two hundred programmers worldwide are now involved in the project.

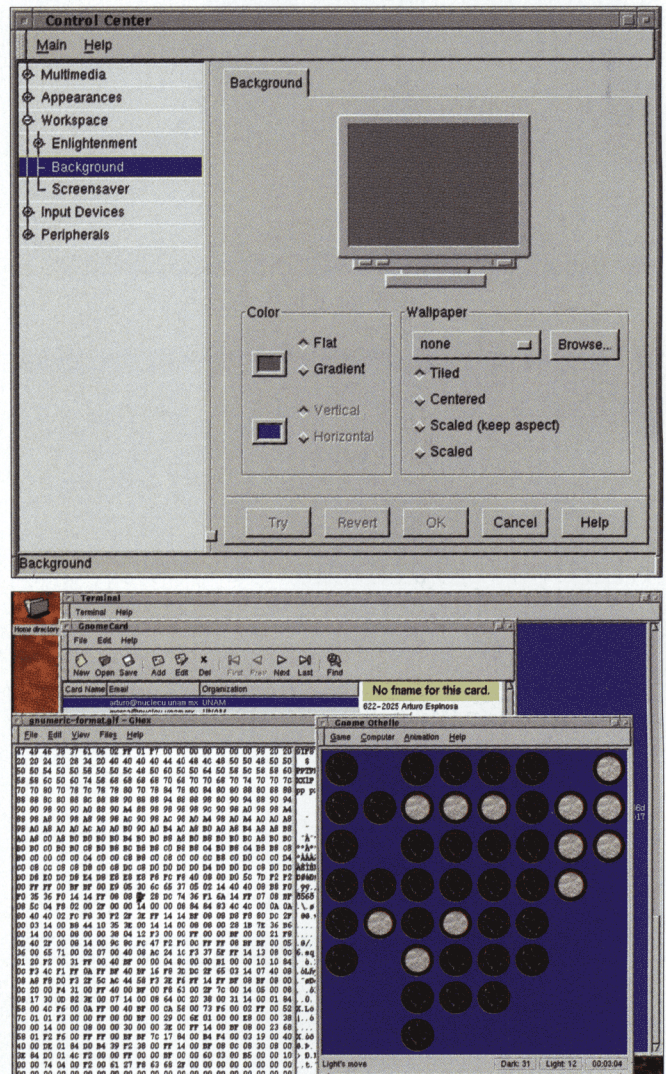
The original announcement called for developers to shape the GNOME project in a number of forums: the GNU announce mailing lists; the Guile mailing list; and the GTK+ and GIMP mailing lists. The programmers and other people who influenced the project were mainly free software enthusiasts with diverse areas of expertise, including graphics programming and language design.

The GNOME team has been working steadily toward creating a foundation for future free software development. GNOME provides the toolkit and reusable component set to build the free software end users are eager for.

Our recent releases of the GNU Network Object Model Environment have been GNOME 0.20, the first version of GNOME that showed signs of integrations, released in May 1998; the Drooling Macaque 0.25 release, with more features; and finally, our latest public release, GNOME 0.30, code named Bouncing Bonobo.

The GNOME 0.20 release was the first release included in a CD-ROM distribution. Red Hat 5.1 shipped with a technology preview of the GNOME desktop environment and it was first demonstrated at the 1998 Linux Expo in North Carolina.

Before the Drooling Macaque release, GNOME software releases were coordinated by two or three people on the team. This became a significant burden, as precious time was being used coordinating each release. We have been trying to make the release process more modular and have assigned different modules to package maintainers. Each package maintainer is responsible for packing, testing and releasing their packages independently of the main distribution, which we consider to be the core libraries and the core desktop applications. So far we have had some success, but there is still room for improvement. We will continue to polish the release process to make it simpler.



The most recent GNOME release, Bouncing Bonobo, is the first to feature the GNOME spreadsheet Gnumeric.

Red Hat Advanced Development Labs

In January 1998, Red Hat announced the creation of the Red Hat Advanced Development Laboratories (RHAD). The initial objective of Red Hat Labs was to help the GNOME effort by providing code and programmers and by helping us manage the project resources.

All code contributed to GNOME by Red Hat Advanced Laboratories has been provided under the terms of the GNU GPL and the GNU LGPL licenses. Several GTK+ and GNOME developers have been hired by Red Hat and they have rapidly provided the GNOME project with a number of important features.

For example, Rasterman has implemented themes for GTK+; the GTK+ themes allow a user to change the appearance of the widgets. This is done by abstracting the widget drawing routines from the toolkit and putting those drawing

The Best Things Come In Small Packages

AlphaPoweredTM

The Cobra RX SlimLine is an Alpha parallel processing node which delivers the most computing power per square inch in the industry. The custom rackmount chassis is only 2U high (3.5 in.), allowing for a densely packed Linux cluster consisting of up to 24 Alpha nodes in an 84" rackmount cabinet. That equates to 12.7-16 GigaHertz (or a theoretical 3.3-6.7 gflops) of pure power whenever and wherever you need it. With our expertise and service, it's no wonder Carrera Alpha Solutions were chosen to power Digital Domain's "Titanic" Linux "render ranch" and Beowulf Alpha Linux clusters such as the Los Alamos Laboratories/Avalon project. Whether you need a single node or a networked array, Carrera has a blazing solution for your Linux application at a price that fits your budget.



Cobra RX SlimLine

Cobra RX Cluster

Cobra RX 4200: Up to 12 Nodes/42" cabinet
Cobra RX 6300: Up to 18 Nodes/63" cabinet
Cobra RX 8400: Up to 24 Nodes/84" cabinet

**Call Today to create your custom
Linux solution.**

1-800-576-RISC
www.carrera.com



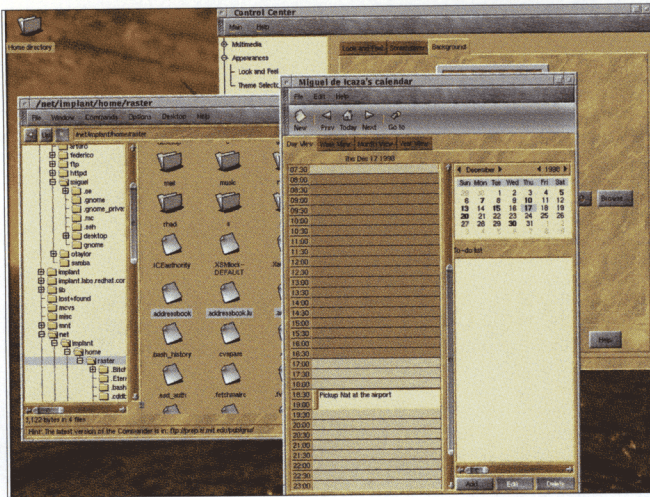
Cobra RX SlimLine

Dimensions: Standard 19" Rack
Node Options:

| | |
|-------------------|------------------------------------|
| CPU: | Alpha 21164@ 533 - 667 Mhz |
| Cache: | 1MB L2, 2 - 4MB L3 |
| Memory: | 128 - 1024 SDRAM (x72 DIMMS) |
| Bus Speed: | 66 MHz |
| Controller: | IDE, 40MB/Sec Ultra SCSI |
| Lan: | 10/100 Ethernet |
| Hard Drive: | 4.5, 9, or 18 GB |
| Graphics: | 2 MB DRAM, Headless |
| CD-ROM: | 32X IDE (Optional) |
| Floppy Drive: | 3.5" 1.44MB Floppy Drive |
| Case: | ATX 3.5" (2U) Custom 19" Rackmount |
| Operating System: | Windows NT, LINUX |
| Warranty: | 24 Month Limited Warranty |

COOL
it works with
LINUX

Carrera



routines in modules that can be loaded at runtime. Thus, the user can change the appearance of applications without shutting them down or restarting the desktop.

GTK+ themes are fully working. So far, a number of theme front-ends have been written. At the time of this writing, available themes include Motif, Windows95, Metal, native-GTK+ and a general purpose Bitmap-based engine (see Resources). The web site <http://gtk.themes.org/> keeps an up-to-date list with many contributed themes from which to choose.

Various important changes to the GTK+ toolkit required for the GNOME project, such as the menu keyboard navigation code and the enhanced "Drag and Drop" protocols (XDND and Motif DND), were written by Owen Taylor, a famous GTK+ hacker now working for Red Hat Labs.

Assorted applications were created or are maintained nowadays by the GNOME team at RHAD as well: the Ghostscript front end (by Jonathan Blandford), the GNOME Help Browser and the GNOME RPM interface (Marc Ewing and Michael Fullbright), the GNOME Calendar and GNOME Canvas (Federico Mena) and the ORBit CORBA 2.2 implementation (Elliot Lee).

Other Donations

The GNOME project received a monetary donation from the GNU/Linux Debian team in the early stages of the project, as well as an Alpha board from Quant-X Service and Consulting GmbH. We are very grateful for their contributions.

Some Key GNOME Features

The GNOME libraries provide a framework to create consistent applications and to simplify the programmer's task. More features of the GNOME libraries will be described later. Some of the most important current developments in the GNOME libraries are discussed here.

Metadata

One problem faced in a desktop environment is the fact that it is usually necessary to have a mechanism for storing information about a file's properties. For example, applications might want to bind an icon for a specific executable file, or bind a small thumbnail image for a graphic produced by a graphics program. These icons should be semantically attached to the main file.

The Macintosh OS, for example, provides a way to store this information in the file as its "resource fork". This mechanism would be awkward at best to implement in a UNIX envi-

ronment. The main problem is that a non-metadata-aware application can cause the metadata information to get out of sync.

The GNOME metadata was implemented by Tom Tromey at Cygnus, given a number of design constraints and tradeoffs (described in detail on their web site). The following is a list of the GNOME metadata features:

1. Binding the information on a per-file basis in a per-user setting, and each user keeps track of his own bindings. System defaults apply on top of these.
2. Binding information by file content is done according to the file type using file signatures, similar to the UNIX **file** command.
3. Binding information by a regular expression: for example, a default icon for gif files would be provided by the regular expression `*\.*gif$`.
4. The metadata system is optimized to provide a coherent GUI solution, rather than as a compromise or kludge to existing command-line tools.
5. Most ordinary uses of files will continue to work without metadata, just as they do now.

A number of standard properties for file metadata are available in GNOME. For example, "View" stores the action for viewing the file contents; "Open" stores analogous action for editing; "Icon", which contains the icon, is used for displaying the file on the desktop.

Metadata types are MIME types.

Canvas

GNOME provides a **Canvas** widget, patterned after Tk's excellent canvas. This widget simplifies the programming of applications that need control over graphical components. The most noticeable feature of GNOME Canvas is that it provides a flicker-free drawing area where high-level objects can be inserted and manipulated. Basic zoom and scroll facilities are also a part of the canvas.

The high-level objects inserted into the canvas behave like regular widgets. They can receive X events, grab the focus and grab the mouse just like a regular widget. As with their Tk counterparts, GNOME Canvas items can have their properties changed at runtime with a Tk-like configuration mechanism.

The GNOME Canvas ships with a number of items derived from the

GnomeCanvasItem object: lines, rectangles, ellipses, arrows, polylines and a generic widget container to embed GTK+ widgets within a canvas. The Canvas framework is designed to be very extensible. As proof of this extensibility, the GNOME spreadsheet is implemented on top of the base canvas engine, with additional functionality provided by spreadsheet-specific CanvasItems.

Note that the current Canvas uses Gdk primitives (a thin wrapper over Xlib primitives) to draw, so it is limited in the quality and range of special effects that can be provided with it, which bring us to the next step in Canvas technology.

66

LINUX JOURNAL FEBRUARY 1999

Raph Levien is working on an advanced rendering engine for the Canvas. It was originally developed as a stand-alone widget within his Type1 outline font editor, **gfonted**. At the time of this writing, work on integrating the engine into the Canvas was getting underway.

Features of this engine include:

- Anti-aliased rendering of all items
- Alpha transparency
- Items for vector and bezier paths
- Items for RGB and RGB plus alpha images
- Vector operations, including clip (intersect), union, difference and stroke layout
- PostScript Type1 font loading and rendering

The engine's design goal is to support almost all of the PostScript imaging model with the addition of alpha transparency. As such, it is expected to be an excellent starting point for high-powered graphics applications.

In spite of the ambitious goal of keeping the display up to date with entirely anti-aliased and alpha-composited items, performance is surprisingly good—comparable, in fact, to the Xlib-primitive-based canvas engine.

His code is expected to be merged into the main Canvas sometime soon.

Window Manager Independence

GNOME does not depend on a special window manager—any existing window manager will do. GNOME specifies window manager hints that can be implemented by the window manager to give the user better desktop integration, but they are optional. The E window manager implements all of the GNOME window manager hints and can be used as a reference implementation for people wishing to extend their window managers to be GNOME-compliant. The ICEWM manager is tracking those developments and is also considered to be a GNOME-compliant window manager, although at this time, it is lagging a bit behind. A few people have shown interest in providing the WindowMaker and FVWM2 maintainers with patches to make those window managers GNOME-aware.

Component Programming

Historically, one of the attractions of UNIX has been the philosophy of

small tools that each do one thing well and combining these tools, using pipes and simple shell scripts, to perform more complex tasks. This philosophy works very well when data objects are represented as plaintext and operations are effectively filters. However, this UNIX command-line philosophy does not scale well to today's world of multimedia objects.

Thus, it would be nice to have a framework in GNOME that would provide software reuse and component plugging and interaction, i.e., connecting small specialized tools to carry out complex tasks. With this infrastructure in place, GNOME applications can once again return to the UNIX roots of simple, well-specialized tools.

An RPC system was then required for providing this sort of functionality, so we decided to use CORBA (the Common Object Request Broker Architecture) from the Object Management Group (OMG). CORBA can be thought of as an object-oriented RPC system, which happens to have standardized bindings for different languages.

CORBA opened a range of applications for us. Component programming allowed us to package programs and shared libraries as program servers that each implement a specific interface.

For example, the GNOME mail program, Balsa, implements the `GNOME::MailMessage` interface that enables any CORBA-aware pro-

Your Code is Tight, Fast and Elegantly Simple.

It Deserves A Database Built the Same Way.



For years, serious C and C++ programmers have relied on Raima's proven database technology for the ultimate in fast and efficient code. Now Java™, Delphi™ and Visual Basic™ programmers can utilize this same technology in their applications. So if you need:

- ▶ A thoroughly optimized database engine
- ▶ Intuitive and powerful APIs
- ▶ Efficient database architecture
- ▶ Cost efficiency
- ▶ Portability
- ▶ Low overhead/small footprint
- ▶ Blazing data access

Check out our database technology.

Visit our Web Site at:
www.raima.com
or call 1-800-327-2462

If it's fast—it's
RAIMA™

gram to remotely compose and customize the contents of a mail message and send it. Thus, it is possible to replace the mail program with any program that implements the `GNOME::MailMessage` interface. As far as the GNOME desktop is concerned, the process just implements the `GNOME::MailMessage` interface. This means, for example, that I will be able to continue using GNUS to read my mail and have GNUS completely integrated with the rest of my desktop. This also applies to the other components in the GNOME system: the address book, file manager, terminal emulation program, help browser, office applications and more.

Beside providing the basic GNOME interfaces, applications can provide an interface to their implementation-specific features. This is done by using CORBA's interface inheritance. A specific interface would be derived from the more general interface. For example, GNUS would implement the `GNOME::MailMessage` interface and extend it with GNUS-specific features in the `GNOME::GnusMailMessage` interface. This interface would hypothetically allow the user to customize GNUS at the Lisp level, something other mailers may not do. Another example would be a `GNOME::MozillaMailMessage` interface that would let a user configure the HTML rendering engine in Mozilla mail.

Not only does CORBA address these issues, but it can also be used as a general interprocess communication engine. Instead of inventing a new ad-hoc interprocess communication system each time two programs need to communicate, a CORBA interface can be used.

Embedding documents into other documents has been popularized by Microsoft with their Object Linking and

Embedding architecture. A document-embedding model similar in spirit is being designed for GNOME (the Baboon model), and all of the interprocess communication in this model is defined in terms of CORBA interfaces.

Initially, we were very excited by the possibilities presented by CORBA, but we soon realized that using CORBA in the GNOME desktop was going to be more difficult than we expected.

We tried using Xerox's ILU for our CORBA needs. The license at the time did not permit us to make modifications to the code and redistribute them, an important thing for the free software community, so we had to look for alternatives. Xerox has since changed the licensing policy.

After evaluating various free CORBA implementations, we settled on MICO, as it was the most feature-full free implementation. MICO was designed as a teaching tool for CORBA, with a primary focus on code clarity.

Unfortunately, we soon found that MICO was not a production-quality tool suitable for the needs of GNOME. For one, we found that the rather indiscriminate use of C++ templates (both in MICO and in MICO-generated stubs) proved to be a resource hog. Compiling bits of GNOME required as much as 48MB of RAM for even the simplest uses of CORBA, and this was slowing down our development. Another problem was that MICO supported only the C++ CORBA bindings. Even though an initial attempt had been made at providing C bindings, they were incomplete and not well-maintained.

To address these problems, Dick Porter at i2it and Elliot Lee at Red Hat Labs wrote a C-based, thin and fast CORBA 2.2 implementation called ORBit. As soon as ORBit became stable, the use of CORBA throughout GNOME began, after a delay of almost eight months.

With an efficient, production quality CORBA implementation under our control, we can ensure that CORBA-enabled interprocess communication is a valuable service for application programmers, rather than a source of overhead and bulk.

Dissecting a GNOME Desktop Application

The toolkit

GNOME desktop applications have been built on top of the object-oriented GTK+ toolkit originally designed as a GUI toolkit for the GNU Image Manipulation Program (GIMP).


GTK+ has been implemented on top of a simple window and drawing API called Gdk (GTK Drawing Kit). The initial version of Gdk was a fairly thin wrapper around the Xlib libraries, but a port to Win32 and a port to the Y windowing system are presently in alpha stages.

GTK+ implements an object system entirely in C. This object system is quite rich in functionality, including classical single inheritance, dynamic creation of new methods and classes, and a "signal" mechanism for dynamically attaching handlers to the various events that occur in the user interface. One of GTK's great strengths is the availability of a wide range of language bindings, including C++, Objective-C, Perl, Python, Scheme and Tom. These language bindings provide access both to GTK+ objects and to new objects programmed in the language of choice.

An additional feature of GNOME is Rastman's Imlib library. This library is implemented alongside Gdk and provides a fast yet flexible interface for loading and saving images and rendering them on the screen. Applications using Imlib have quick and direct access to PNG, GIF, TIFF, JPEG and

Get LINUX from the /root.

- ◆ Linux Distributions
- ◆ Linux Software
- ◆ Official Red Hat Products
- ◆ Manuals from the FSF
- ◆ Serial Display Modules
- ◆ and more...




linux central
the /root for linux resources

http://linuxcentral.com/

The World's fastest computers use PGI Compilers and Tools. What are you using?

When HPC users need production-quality parallel compilers and tools, they turn to The Portland Group, Inc. (PGI). In fact, PGI's *PGHPF*® High Performance Fortran compiler is in use on nearly 40% of the world's TOP 100 computers. No other parallel Fortran compiler even comes close. Not SGI. Not IBM. Not Sun. And now PGI is leveraging that same power onto your desktop with a complete line of parallel Fortran, C, and C++ compilers and tools for Linux and Windows® NT workstations and servers. Whether you're developing applications for your workstation, or for the world's fastest computers, PGI compilers will harness the full power of your system. Look us up on the web, and let PGI help you get your parallel applications in production.

It's not a Workstation without PGI Compilers and Tools.



The Portland Group

9150 SW Pioneer Ct. Suite H Wilsonville, OR 97070
Phone: (503) 682-2806 FAX: (503) 682-2637 e-mail: sales@pgroup.com
<http://www.pgroup.com>

XPM files as well as other formats available through external conversion filters.

The Support Libraries

C-based GNOME applications use the glib utility library. **glib** provides the C programmer with a set of useful data structures: linked lists, doubly linked lists, hash tables (one-to-one maps), trees, string manipulation, memory-chunk reuse, debugging macros, assertion and logging facilities. **glib** also includes a portable interface for a dynamic module facility.

The GNOME libraries

The GNOME libraries add the missing pieces to the toolkit to create full applications, dictate some policy, and help in the process of providing consistent user interfaces as well as localizing the GNOME applications so they can be used in various countries.

The current GNOME libraries are GTK+-Xmhtml, gnome-print, libgnome, libgnomeui, libgnorba, libgtop, gnome-dom and gnome-xml. Other libraries are used for specific applications: libPropList (soon to be replaced by a new configuration engine) and audiofile.

The main non-graphical library is called libgnome. This provides functions to keep track of recently used documents, configuration information, metadata handling (see below), game score functions and command-line argument handling. This library does not depend on the use of a windowing system.

As we use CORBA to achieve parts of our desktop integration, we have a special library called the libgnorba library to deal with various CORBA issues. It provides GUI/CORBA integration (to let our GUI applications act as servers), authentication within the GNOME framework and service activation.

The gnomeui library, on the other hand, has the code that requires a window system to run. It contains the following components:

- The GNOME session management support
- Widgets, both as straightforward extensions of GTK+ and designed to be dependent on libgnome features
- A set of standard dialog boxes otherwise not available on GTK+, well-integrated with other GNOME libraries
- Standard property configuration dialog boxes
- Standard top-level window handling
- A multi-document interface (gnome-mdi)
- Windowing hints
- CORBA integration where required

GTK+-XmHTML is a port of the Koen D'Hondt's XmHTML widget for Motif and is used for our HTML display needs. Our changes are being folded back into the main distribution.

The libgtop library allows system applications to be easily ported to various operating systems; it provides system, process and file system information.

gnome-xml provides XML file loading, parsing and saving for GNOME applications and is being used in the GNOME spreadsheet (**Gnumeric**) and in the GNOME word processor program. **gnome-dom** provides an implementation of the World Wide Web Consortium's Document Object Model for GNOME applications. By the time you read this article,

gnome-dom will have been deployed widely in the GNOME office applications. Both gnome-xml and gnome-dom were developed by Daniel Veillard from the World Wide Web Consortium.

gnome-print implements GNOME's printing architecture. It consists of a pluggable rendering engine as well as a set of widgets and standard dialog boxes for selecting and configuring printers. In addition, gnome-print is responsible for managing outline fonts and contains scripts that automatically find fonts already installed on the system.

The GNOME print imaging model is modeled after PostScript. Basic operations include vector and bezier path construction, stroking, filling, clipping, text (using Type1 fonts, with TrueType to follow shortly) and images.

Currently, gnome-print generates only PostScript output. However, the design of the imaging model is closely synchronized with the anti-aliased rendering engine for the Canvas and it is expected that these two modules will be interoperating soon. In particular, it will be possible to "print" into a canvas (useful for providing a high-quality screen preview) and to print the contents of a canvas. This feature should simplify the design of applications which use the Canvas, as very little extra code will be needed to support printing.

The same rendering engine will be used to render printed pages directly without going through a PostScript step. This path is especially exciting for providing high-quality, high-performance printing to color ink-jet printers, even of complex pages containing transparency, gradients and other elements considered "tricky" in the traditional PostScript imaging model.

Bindings

One explicit goal of GNOME was to support development in a wide range of languages, because no single language is ideal for every application. To this end, bindings for both GTK+ and the GNOME libraries exist for many popular programming languages, currently C, C++, Objective-C, Perl, Python, Scheme and Tom.

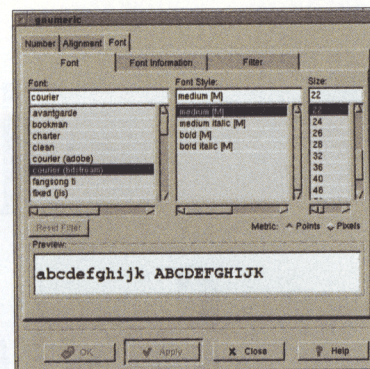
The early involvement of Scheme, Tom and Perl hackers in both the GTK+ and GNOME projects has helped in making the GTK+ and GNOME APIs easy to wrap up for various languages. Multi-language support is "baked in" to the design of GTK+ and GNOME, rather than being added on as an afterthought.

Development model

GNOME is developed by a loosely coupled team of programmers around the world. Project coordination is done on the various GNOME mailing lists.

The GNOME source code is kept on the GNOME CVS server ([cvs:cvs.gnome.org/cvs/gnome/](http://cvs.cvs.gnome.org/cvs/gnome/)). Access to the source code through Netscape's Bonsai and LXR tools is provided at <http://cvs.gnome.org/>, to help programmers get acquainted with the GNOME source code base.

Most developers who have contributed code, major bug fixes and documentation to GNOME have CVS write access, fostering a very open atmosphere. GNOME developers come from a wide range of backgrounds and have diverse levels of



skill and experience. Contributions from less experienced people have been surprisingly helpful, while the older, wiser coders have been happy to mentor younger contributors on the team. The GNOME developer community values clean, maintainable code. Even programmers with many years of coding experience have noted how involvement with the GNOME project has helped them write better code.

The GNOME Office Suite Applications

As the GNOME foundation libraries become more stable, the development of larger programming projects has become possible and has allowed small teams of developers to put together the applications which will make up the GNOME office suite.

As with other GNOME components, the GNOME office suite is currently catching up with commercial offerings. By providing an office suite which is solid, fast and component-based, the code written for the GNOME project might become the foundation for a new era of free software program development.

The office suite leverages a lot of knowledge many of us have acquired during the past year while developing various GNOME components. Our coding standards are higher, the quality is better and the code is more clean and more robust.

The availability of these applications has provided us with the test bed we required to complete our document embedding interfaces (the Baboon model).

Acknowledgments

There is no way to thank all of the contributors to the GNOME project in this space. All of these contributions are gratefully appreciated.

I would like to especially thank Alan Cox, Nat Friedman, Raph Levien and Richard Stallman for reviewing the draft of this document.

Resources

bonobo: The GNOME team has learned that the Bonobo, the primate closest to humans, is an endangered species. If you want to know more about how you can help save the Bonobos, check this web page: <http://www.gsu.edu/~wwwbpf/bpf/>

GIMP: <http://www.gimp.org/>

GNU: <http://www.gnu.org/>

GTK+: <http://www.gtk.org/>

GNOME: <http://www.gnome.org/>

Gnumeric: <http://www.gnome.org/gnumeric/>

gnome-print: <http://www.levien.com/gnome/print-arch.html>

GWP: <http://www.hungry.com/products/gwp/>

OMG: <http://www.omg.org/>

ORBit: <http://www.labs.redhat.com/orbit/>

RHAD: <http://www.labs.redhat.com/>

Themes: <http://www.labs.redhat.com/themes/>

Tom Tromey:

<http://www.cygnus.com/~tromey/gnome/metadata.html>

Y: <http://www.hungry.com/>

Two word processing projects are going on for GNOME: one of them is GWP by Seth Alves at the Hungry Programmers and the other one is Go from Chris Lahey. GWP is currently more advanced and has printing working with the GNOME printing architecture.

Gnumeric, the GNOME spreadsheet project, is aimed at providing a commercial-quality spreadsheet with advanced features. It provides a comfortable and powerful user interface. As with other components in GNOME, we have worked toward providing a solid and extensible framework for future development.

Recently, work has begun on Achtung, the GNOME presentations program. It is still in the early stages of development.

Getting GNOME

Tested source code releases of GNOME are available from GNOME's ftp site: <ftp://ftp.gnome.org/>.


It is also possible to get the very latest GNOME developments from the anonymous CVS servers. Check the GNOME web page for details on how to pull the latest version straight from the CVS servers.

Breaking news about GNOME is posted to the GNOME web site in <http://www.gnome.org/>, along with documents to get you started on GNOME and developing GNOME applications.

Miguel de Icaza is one of the GNU Midnight Commander authors as well as a developer of GNOME. He also worked on the Linux/SPARC kernel port. He can be reached via e-mail at miguel@gnu.ai.mit.edu.

Why reinvent the wheel?

Save time, money & minimize your programming headaches.



NAG Software Productivity Tools.

Numerical Libraries

- C
- Fortran
- Parallel

Compilers

Visualization
Simulation

Numerical Algorithms Group, Inc.

Phone 630-971-2337 FAX 630-971-2706

EMAIL info-lj@nag.com

WEB <http://www.nag.com>

NAG

Virtual Network Computing

Mr. Harvey tells us about the VNC software package and how to set it up to control MS Windows servers from Linux.

by Brian Harvey

In today's changing world, an increasing number of UNIX system administrators are finding they need to support Windows NT servers in their work environments. Whether Exchange or application servers, NT servers are starting to creep into what were once UNIX-only shops. The responsibility of managing an NT server can be discouraging to any UNIX guru. UNIX users are used to the flexibility of the X Window System—the ability to run applications easily on any UNIX server and have remote X applications display on the local desktop. It is much more difficult to manage NT servers remotely and the administrator usually needs to be at the system's console to run most NT applications.

Several commercial products allow MS Windows applications to be controlled remotely from an X desktop. In addition, there are several commercial X servers for MS Windows which allow the opposite. However, until recently an equivalent free software package was not available.

VNC

Researchers at the Olivetti & Oracle Research Laboratory (ORL) have released the VNC software package under the GNU general public license. VNC, which stands for Virtual Network Computing, is a client/server-based, stateless, platform-independent protocol developed at ORL. This protocol implements a remote display system in which a user is allowed to control a computing “desktop” managed by a VNC server, by connecting to it from a VNC client application called a “viewer”. VNC servers currently exist for Windows 95/NT, Macintosh and UNIX. A variety of VNC clients exist as well for a number of operating systems. Figure 1 shows all the connections currently possible with the VNC protocol. Many of the VNC viewers are ported by users on the Net. ORL supplies precompiled server and client binaries for Windows 95/NT, Macintosh, Linux, Digital UNIX and Solaris. In addition, ORL provides a Windows CE client.

In this article, I will discuss how to set up the VNC software, allowing you to control a Windows desktop from Linux running the X Window System (probably the most common use of VNC for Linux users).

Installation

Linux (VNC Viewer)

ORL provides an x86 Linux 2.0 binary that works great with Red Hat 5.1 and can be retrieved from their download page (see Resources). Once you have the package, unarchive it using **gunzip** and **tar**. The binary distribution provides no installation script, but for our purposes we simply need to have root copy the viewer binary, **vncviewer**, into a



suitable location accessible by others, such as /usr/local/bin.

Windows 95/NT (VNC Server)

ORL provides a pre-compiled Windows 95/NT binary supplied as a package that can also be downloaded from their download page. The package installs like most other Windows software packages, i.e., using InstallShield. The VNC server (WinVNC) can be installed as a regular application (started/stopped by the user currently logged on to the console) or as an NT service (starts automatically when NT boots; does not exit when user logs out). Installation as a service is a new feature in recent versions of WinVNC. The latest version at the time of writing, is 3.3.2R5. VNC is actively developed, so a newer version of the software will most likely be ready to download by the time you

read this. I recommend installing WinVNC as a service so that the VNC server is always running and you do not have to remain logged in on the Windows console at all times.

To install WinVNC as a service, simply install the package as you would normally install any other Windows application, then type in a command window:

```
cd PATH_to_installed_WinVNC
WinVNC.exe -install
WinVNC.exe -run # or reboot NT to have the
# service start automatically
```

Configuration

The Linux VNC viewer requires no configuration to use. However, the Windows VNC server does require some minor configuration. To bring up the configuration window either right-click on the WinVNC icon in the Windows




Figure 1. VNC Protocol Connections

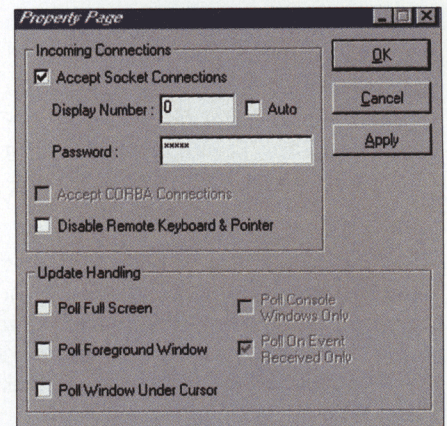


Figure 2. Windows VNC Configuration Window



THE ***MOVERS*** ARE ***HERE.***

Now you can run UNIX® system applications on Windows NT.™ Safely. Within days. Without staffing up. And without hacking up your code.

Just add INTERIX™ to your Windows NT system.

INTERIX – and only INTERIX – puts a complete, proven, high-performance open systems environment directly on top of the Windows NT micro-kernel.

This allows UNIX systems apps to run precisely as written – only now, on Windows NT.

And you can cut months, even years,

per project. Case in point: The Apache server was deployed on Windows NT in a single day.

InfoWorld called it “amazing.” Try it with any application, any size, and you’ll agree.

Any questions? Let us answer them. Want references? We’ve got hundreds.

Each one will attest that INTERIX makes running UNIX system apps on Windows NT more than just possible.

INTERIX makes it a breeze.



**For a no-risk
evaluation copy,**

**as well as
detailed specs
and ordering**

information, click

“info/eval” at

www.interix.com/

INTERIX
SOFTWARE SYSTEMS, INC.



FORMERLY KNOWN AS OPENNT

Interix is a trademark of Softway Systems, Inc. Windows NT is a trademark of Microsoft Corporation. UNIX is a registered trademark of the Open Group.

NT/95 system tray and select **Properties**, or open a DOS command window and type:

```
cd PATH_to_installed_WinVNC
WinVNC.exe -settings
```

In the configuration window, shown in Figure 2, the following options can be set:

- Make sure **Accept Socket Connections** is selected. If this option is not checked, all incoming connections will be disabled.
- The **Display Number** can be left at 0. This value is specified when using a VNC viewer to connect to this server.
- Set a **Password** to secure access to this VNC desktop (a good idea). When connecting to this VNC server via a viewer, you will be prompted for the same password.
- If **Disable Remote Keyboard & Pointer** is selected, all incoming viewer connections will be able to see the desktop but will not be able to move the mouse or type anything (a read-only connection).
- In the **Update Handling** section, various options can be turned on/off to control how the VNC server sends "desktop changes" to a VNC viewer. See <http://www.orl.co.uk/vnc/winvcn.html> for in-depth explanations on the pros and cons of each option.

Press the **Ok** or **Apply** button to apply your configuration changes.

Using the Linux VNC viewer

Once you have WinVNC running on a Windows server, try connecting to it from your Linux desktop by typing (within X) the following command, followed by the password you gave when configuring WinVNC (if any):

```
> vncviewer windows_server:0
vncviewer: VNC server supports protocol version 3.3
(viewer 3.3)
Password: your_password
vncviewer: VNC authentication succeeded
vncviewer: Desktop name "boxster"
vncviewer: Connected to VNC server, using protocol
version 3.3
vncviewer: VNC server default format:
16 bits per pixel.
Least significant byte first in each pixel.
True color: max red 31 green 63 blue 31
shift red 11 green 5 blue 0
Using default colormap and translating to BGR233
Creating window depth 8, visualid 0x22 colormap 0x21
```

If you typed the password correctly, several lines of information will appear and a new large window will pop up showing the entire remote Windows desktop. When you are finished using the VNC viewer, simply close the viewer's window to close the connection. The remote Windows desktop will be left in the last state the viewer left it in.

Figure 3 shows a sample Linux desktop with a newly opened VNC viewer connection "viewing" a Windows NT desktop.




Figure 3. Linux Desktop Viewing Windows NT

A nice feature available in recent VNC releases is the ability to send the infamous **CTRL-ALT-DEL** key sequence to the Windows desktop shown in a VNC viewer. This feature has distinct advantages when the VNC server is installed as a service:

- If the VNC server is installed as a service under Windows NT, you don't need to have a user logged on all the time with the VNC server running as a Windows application. When it comes time to use that server remotely, simply connect to it with a VNC viewer, press **CTRL-ALT-DEL** to get the NT login Window, and log on as you normally would to the NT box.
- If you need to stay logged on to the NT server but want to exit your local X session, you can type **CTRL-ALT-DEL** to get the "Windows NT Security" pop-up window, click on "Lock Workstation" to lock the console, close the VNC viewer connection, then exit your X session. You will still remain logged on to the NT server; its screen is now locked.

Advantages

The VNC protocol has several advantages. The main one is that it is stateless. A user can close a connection to a remote desktop from one VNC viewer and later reconnect to that same remote desktop from the same or different VNC viewer, and it will be in the same state.

When using the Java VNC viewer, a system administrator can control a Windows 95/NT, Macintosh, or UNIX desktop from anywhere in the world using a Java-enabled browser. The VNC server can be configured so that all incoming viewer connections will be able to see the desktop but will not be able to move the mouse or type anything (a read-only connection). This option comes in handy in a teaching environment, where each student in a class connects to the instructor's "desktop" and watches a demonstration on his own computer rather than on an overhead connected to the instructor's computer.

How I Use VNC

At work, I have an Alpha running Digital UNIX and a P133 running Windows NT 4.0.

Although I am strictly a UNIX systems administrator, my company's e-mail standard is based on Microsoft Exchange.

Therefore, I am required to have a Windows desktop on my desk in order to read Exchange e-mail.

However, at home I run only Linux. I was looking for a way to read my Exchange e-mail from home. After reading about VNC, I knew I had found what I was looking for.

I use the Linux VNC viewer at home to connect to the Windows NT box on my desk at work over a PPP connection. Figure 4 shows me reading my Exchange e-mail with such a setup. While VNC performance over a PPP line isn't spectacular, it is very usable and solves my problem of not being able to read Exchange e-mail from home.



Brian Harvey is currently a UNIX Systems Administrator for U.S. Technical Services in Huntington Beach, CA. He is a graduate of UC Riverside with a BS and MS in Computer Science. He can be reached via email at brian.harvey@ustsvs.com

www.uk.research.att.com/vnc/

Resources

Main VNC home page at ORL:
<http://www.orl.co.uk/vnc/>

On-line documentation page for VNC servers and viewers: <http://www.orl.co.uk/vnc/docs.html>

VNC download page for precompiled binaries, source and documentation:
<http://www.orl.co.uk/vnc/download.htm>

Various contributed VNC viewer ports plus hints on how to get the viewer compiled and working on your favorite OS for which ORL does not supply a binary:
<http://www.orl.co.uk/vnc/contribs.html>

To subscribe to the VNC mailing list, send e-mail to vnc-list-request@orl.co.uk with the text "subscribe vnc-list" in the body of the message.

Searchable Linux manual pages:
<http://www.linuxjournal.com/man.html>

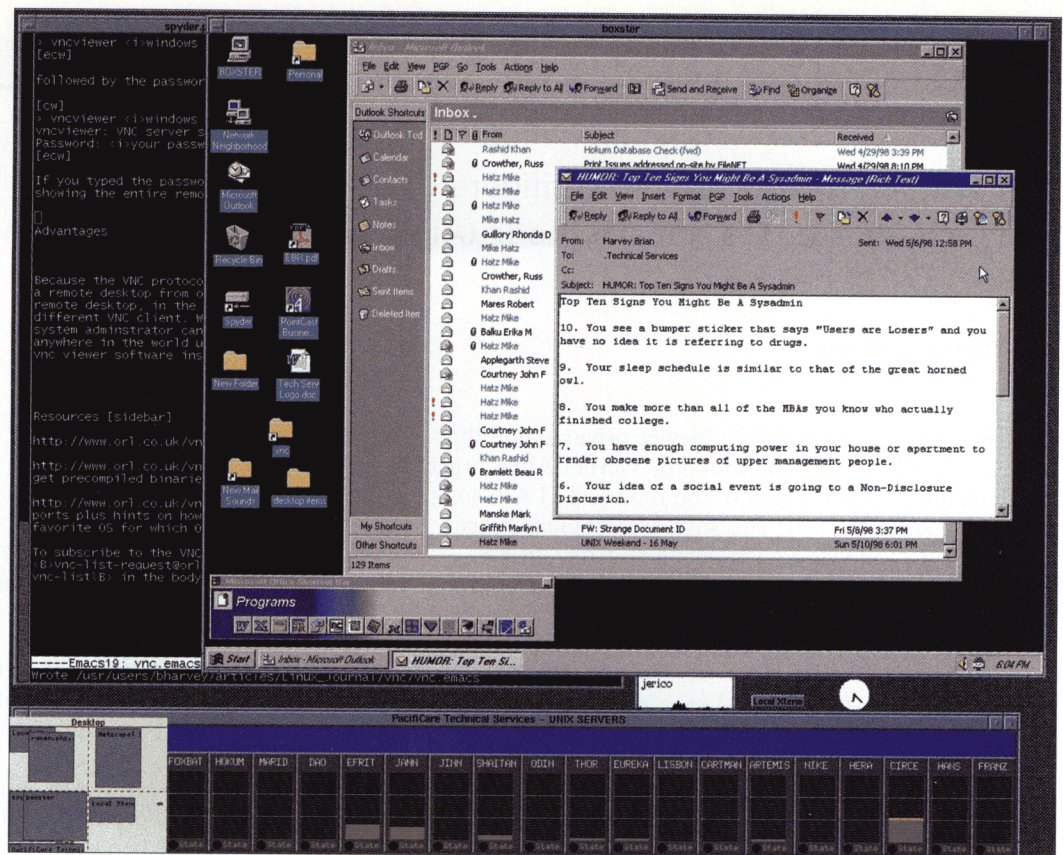


Figure 4. Reading MS Exchange from Linux

VITAL CRISP

WWW.VITAL.COM

WWW.CRISP.COM

The universal editing environment for you.

- Brief, Vi, Emacs, EDT, Wordstar, Windows Emulation.
- Dynamic Class Browser, Tagging.
- 100% Brief Compatible Macro Language.
- Colored Editing & Printing.
- Code Folding & Outlining.
- Dynamic Spell Checking.
- Includes X, Character On Unix / Visual, DOS On Windows At No Additional Cost.

**NEW
RELEASE
6.0
VERSION**

Platforms: Sun, HP, AIX, SGI, Alpha, SCO, Linux, Lynx, Univel, BSD, Win95, NT.



Tel: (972) 612-2684
Fax: (972) 379-6439
Orders: 800-OK-CRISP

FREE
Home Copy
With
Maintenance

Configuring ATM Networks

This article describes how to configure Linux-based PCs and an asynchronous transfer mode (ATM) switch to build an ATM network.

by Wayne J. Salamon

The Linux ATM software (device driver and utilities) is developed and supported by Werner Almsberger in Switzerland as part of the Linux-ATM API software set (see Resources). This software contains device drivers for the following ATM adapters: Efficient ENI-155P, SMC ATM Power 155, Rolf Fiedler's TNETA1570 board, Zeitnet ZN1221/ZN1225 and the IDT 77901/77903 155 and 25 Mbps adapters. Also, a driver for the Fore PCA-200E ATM adapter is available separately (see Resources). The two adapters I have experience with are the Efficient ENI-155p and the Fore PCA-200E.

The National Institute of Standards and Technology (NIST) uses ATM and Fast-Ethernet networks as interconnects in its scalable cluster computing initiative. One research area is evaluating the benefits of ATM and Fast-Ethernet networks in this cluster environment.

In this article, I will tell you how to obtain and install the ATM support software and device drivers. I will also describe how to configure the ATM connections on the PCs and the switch to be used for IP network traffic.

The ATM interface cards I use are ENI-155P ATM adapters produced by Efficient Networks and PCA-200EPC adapters from Fore Systems. These cards are installed in standard Pentium or Pentium-Pro-based PCs running Linux. The ATM switch I used for this article is a Fore ASX-1000, although the information I give applies to all of the Fore ATM switches. This switch can be set up to allow the Linux workstations to use IP over both Switched Virtual Circuits (SVC) and Permanent Virtual Circuits (PVC).

Obtaining and Installing the Linux-ATM Software

The ATM software is available from <http://lrcwww.epfl.ch/linux-atm/>. The software is packaged as a compressed, gzipped tar file. Each version of the software is tied to a specific version of the Linux kernel. For this article, I used version 0.35 running on Linux kernel 2.1.90. The size of the ATM software distribution is roughly 500KB. The device driver for the Fore PCA-200E adapter can be obtained by anonymous FTP from <ftp://os.inf.tu-dresden.de/pub/pca200e/>. Refer to the README file in the PCA200 distribution for further information.

The driver portion of the Linux-ATM software, as well as the changes to the Linux kernel, are shipped as one large patch file. Therefore, adding support to the Linux kernel for ATM is straightforward: apply the kernel patch, configure and rebuild the kernel in the usual way. The ATM configuration items you must have are:

- Asynchronous Transfer Mode (ATM) (CONFIG_ATM)

- Classical IP over ATM with ATMARP (CONFIG_ATM_ATMARP)
- Device driver, one of the following:
 - Efficient Networks ENI155P (CONFIG_ATM_ENI)
 - ZeitNet ZN1221/ZN1225 (CONFIG_ATM_ZATM)
 - Rolfs TI TNETA1570 (CONFIG_ATM_TNETA1570)
 - IDT 77201 (NICStAR) (CONFIG_ATM_NICSTAR)

I recommend starting with a fresh Linux kernel source tree before applying the ATM patch. Refer to the USAGE file that is part of the Linux-ATM software, as things may change. All of the device drivers in the distribution can be built as kernel modules or as part of the kernel object itself. If you are using a Fore PCA-200E adapter, you do not select a driver during the kernel configuration. The PCA-200E device driver is built as a module separately, as specified in the README file included in the PCA200 distribution.

ATM is a point-to-point, switched technology.

After the kernel is patched, rebuilt and installed, you are ready to build the ATM support software. Again, refer to the instructions in the USAGE file. One change I recommend is installing the support files in `/usr/local/atm-version/bin` and creating a soft link from `/usr/local/atm` to the actual install directory. By using the soft link, you can change ATM software levels and back them out, if needed, without changing the configuration scripts.

Configuring the ATM Device Interface

You are now ready to configure the IP over ATM. First, you must decide what type of "virtual circuits" to use to connect the machines. ATM is a point-to-point, switched technology; in order for two hosts to communicate, a virtual circuit must be established between them.

Switched Virtual Circuits (SVCs) are connections that are established dynamically and torn down when the connection is no longer needed. However, a high latency is associated with establishing a connection. Also, SVCs are deleted after a timeout period if no traffic is sent over the connection. Therefore, the latency associated with SVCs is not always predictable. I encountered several problems when using SVCs, such as connections not being established or sometimes failing to remain open.

Permanent Virtual Circuits (PVCs) are established and kept open. Thus, no latency is associated with establishing the connection, as there is when using SVCs. The disadvantage of PVCs is that the switch must be configured to

establish all the connections between the hosts. When you have several hosts and each host needs to communicate with all the others, the number of PVCs required within the ATM switch grows rapidly. Specific configuration information for SVCs and PVCs is discussed later, but I will jump ahead a bit in order to complete the IP configuration now. The steps to configure the ATM interface are as follows:

- Start the ATM software daemons with these commands:

```
atmsigd -b
ilmid -b
atmarpd -b
```

- Create the ATM device name:

```
atmarp -c atm0
```

- Configure the ATM interface for IP:

```
ifconfig atm0 ipaddr netmask netmask mtu mtu
```

- Add the route for the ATM subnet:

```
route add -net network netmask netmask atm0
```

- Create a permanent ATM ARP (address resolution protocol) cache entry for the ARP server:

```
atmarp -s arpserver arpsrvnsap arpsrv
```

`ipaddr` is the IP address of the ATM interface, `netmask` is the network mask and `network` is the IP address of the network to which we are connecting. `arpserver` is the IP address of the ATM ARP server and `arpsrvnsap` is the ATM address of the ARP server. The ATM ARP server is used to convert an IP address to an ATM network service access point (NSAP) address. (The NSAP address is similar to a media access control (MAC) address and is 20 octets long.) The NSAP address is needed to establish SVCs between nodes. You can also create an `/etc/hosts.atm` file to contain the IP to NSAP mapping, allowing for quicker IP to NSAP translations. For my network, I use the Fore switch as the ARP server. The **atmarpd** daemon maintains a cache of IP to NSAP mappings. The **atmarp** command makes the ARP cache entry permanent when the `arpsrv` option is used.

One final note: if you are going to use PVCs only, you do not need to start the **atmsigd** and **ilmid** daemons. Listing 1 contains a complete example of configuration commands.

Configuration of Switched Virtual Circuits

The ATM switch configuration commands I use apply to the entire family of Fore ATM switches, because they all have a similar command interface.

When using SVCs, a host must pass information to the ATM switch, declaring its intent to set up a connection with another host. The term for connection setup is “signaling”. The ATM protocol used between a host and a switch is the user-network-interface (UNI) signaling standard. There are several revisions of the UNI standard. The Fore ATM switch supports UNI 3.0, 3.1 and 4.0. The Linux-ATM software also supports these versions.

However, there are standards and there are implementations. In setting up our SVCs, I encountered several problems with UNI 3.0 signaling. The UNI 3.1 signaling was more

Listing 1: Script to Configure the ATM Interface

```
#!/bin/sh
#
# Initializes the atm interfaces on
# your system
IPADDR="192.168.1.10"
NETMASK="255.255.255.0"
NETWORK="192.168.1.0"
BROADCAST="192.168.1.255"
ARPSERVER="192.168.1.1"
ARPSRVNSAP="47.0005.80.f21c.0d9f.0020
481c0d9f.00"
#
# NOTE: We change the SDU size below when
# using the Efficient ENI-1* *55P
# adapter cards.
#
SDU_SIZE="4352"
MTU_SIZE="4344"
#
# start up the daemons needed for UNI
# signaling and ATMARP
# (If you are using PVCs only, you don't
# need atmsigd and ilmid, but
# you still need atmarpd for PVCs!)
/usr/local/atm/bin/atmsigd -b
/usr/local/atm/bin/ilmid -b
/usr/local/atm/bin/atmarpd -b
#
# create the atm interface
/usr/local/atm/bin/atmarp -c atm0
# configure the atm interface
/sbin/ifconfig atm0 ${IPADDR} broadcast \
    ${BROADCAST} netmask ${NETMASK} * \
    *mtu ${MTU_SIZE} /sbin/route add -net
${NETWORK} netmask ${NETMASK} atm0
#
# set the atmarp server nsap address
/usr/local/atm/bin/atmarp -s ${ARPSERVER} \
    ${ARPSRVNSAP} qos ubr:sdu=9188 * *arpsrv
#
# tell atmarp the max sdu for the network
# interface
/usr/local/atm/bin/atmarp -q ${NETWORK} \
    ubr:sdu=${SDU_SIZE}
#
# The next line creates a PVC from this
# node to node2 over
# VPI 0, with VCI 70. Note that you can
# use the host name from the /etc/hosts file.
#
/usr/local/atm/bin/atmarp -s node2 0.0.70
```

stable and reliable. To change the signaling on the Fore ATM switch, each port must be changed individually, using the switch control processor (SCP) command interface.

First, log on to the SCP via a TELNET session or by using a terminal attached to the serial port on the Fore switch. The command syntax used here is the same as Fore's. Required parameters are shown between “<” and “>”; options are enclosed in brackets (“[” and “]”); modifiers to options are enclosed in parentheses. One of the modifiers must be chosen.

Change to the UNI configuration menu:

```
localhost::> conf uni
```

The switch prompt is shown in italics, while the command is shown in normal text.

The command **show** will list the current UNI status for each port. If the port is already configured for UNI 3.1, no change needs to be made. Otherwise, you must first delete the current configuration. The syntax for the delete command is `del port vpi`, where *port* is the switch port and *vpi* is the virtual path identifier (usually 0). To delete the signaling on port 1A1 for VPI 0, you would enter this command:

```
localhost::configuration uni> del 1a1 0
```

Now you are ready to configure the port for UNI 3.1. The syntax for the new command is:

```
new <port><vpi> [auto | uni30 | uni31] [-ilmi (up | down)]
```

The **ilmi** option is used when you want the port to respond to integrated local management interface (ILMI) requests. ILMI is used by the hosts to obtain the ATM NSAP address assigned to the host. You usually want to have ILMI active for the port, so the command for port 1A1, VPI 0 is:

```
localhost::configuration uni> new 1a1 0 uni31 -ilmi up
```

Now that the ATM switch ports have been configured, the software on the workstation must be set up. The key portions of the Linux-ATM software are three daemons: **atmsigd** to handle signaling (UNI), **ilmid** to handle ATM address registration and **atmarpd** to map ATM addresses to IP addresses. Listing 1 is the startup script I use to start the ATM daemons and to configure the ATM interface on a host. This script can be called from a system startup script (/etc/rc.d/rc.local, for example) to configure the ATM interface at boot time.

NIST is using ATM and Fast-Ethernet networks as interconnects in its scalable cluster computing initiative.

The ATM signaling daemon, **atmsigd**, must be compiled specifically for the version of signaling you wish to use and must be compatible with the signaling version the ATM switch port has been configured to use. The default version used in the ATM software is UNI 3.0. If you've configured the switch to use UNI 3.1, having the hosts use UNI 3.0 will most likely work, due to backward compatibility. However, I recommend you configure the Linux-ATM software to use the same version as the switch, UNI 3.1.

To have the signaling daemon use UNI 3.1, edit the Rules.make file in the ATM source directory (/usr/src/atm if you follow the steps in the USAGE file). You need to change the STANDARDS line to specify the version of signaling to support. For UNI 3.1, this line should be: `STANDARDS=-DUNI31`.

Special Configuration for ENI-155p ATM Cards

If you are using Efficient ENI-155p ATM cards, the number of simultaneous virtual channels available is limited. The ENI card performs the segmentation and reassembly (SAR) of ATM cells by using memory on the adapter card as a buffer. The host ATM software allocates buffer space for each virtual channel. If you attempt to open

more SVCs than are supported by the available buffer space, you will receive this error message from the ATM ARP daemon:

```
atmarpd:IO: [2]connect: No buffer space available
```

When IP over ATM is used, the device driver sends packets to the ATM card using an ATM Adaptation Layer (AAL). While several adaptation layers are available, AAL-5 is used for IP. The AAL-5 packet is a type of service data unit (SDU) and is somewhat analogous to an Ethernet frame. The AAL-5 packets are divided into individual ATM cells by the Efficient ATM adapter.

The MTU (maximum transmission unit) size for the ATM interface depends on the SDU size. The IP over ATM (Classical IP) specification says that the MTU should be no larger than 9180 bytes. There are also 8 bytes for an AAL-5 trailer, so the SDU for IP over ATM is 9188 bytes in the default configuration. The amount of buffer space needed on the card depends on the maximum SDU size.

The Linux-ATM software allocates three times the maximum SDU size, rounded up to the nearest power of two. In the default configuration, this allocation results in 32KB of buffer space being reserved for each ATM connection ($9180 \times 3 = 27540$, rounded to 32768 bytes). Also, using classical IP causes two SVCs to be made: the initiating machine opens an active connection to the target machine and the target machine opens an active connection back, that is, a passive connection on the initiator. Therefore, these two connections result in the allocation of two buffers on the ATM card, for a total of 64KB.

The default configuration allows a host to have a maximum of fourteen simultaneous connections when using the "client" version of the ENI-155p ATM card, which has 512KB of memory and 504KB of memory available for the SAR buffers. These fourteen connections allow communication using IP over ATM to seven other hosts when using SVCs. If you set up PVCs, you can communicate with fourteen other hosts. When using an ARP server, you have one less connection available, reducing the host count by one as well. The "server" version of the ENI 155p card has 2MB of memory, with 2040KB for SAR buffers, allowing for more simultaneous connections.

To increase the number of simultaneous connections for classical IP, you need to change the size of the maximum SDU set on the ATM interface. By using the allocation rule given above, you can estimate the amount of memory needed for the connections. For example, if you want to use 16KB for each connection, the maximum SDU would be $16384 \div 3$, which is 5461 bytes. I'll use an SDU of 4352 bytes for my example in this article.

The maximum SDU is specified as an option to the ATM ARP daemon. However, when the SDU is changed, the IP interface must also be configured to have an MTU of the same size as the SDU, minus 8 bytes for the AAL-5 trailer. Therefore, in my example the MTU is 4344 bytes.

A potential problem occurs when changing the maximum SDU for the interface: the ATM ARP daemon (**atmarpd**) may not communicate with the ARP server on the Fore switch. Our switch would accept only connections with an SDU of 9188 bytes. The fix for this problem is to create a permanent ARP cache entry on the host, specifying the maximum SDU of 9188 bytes, for the connection to the ARP server. The steps for configuring the ATM software on the workstation are as follows:

- Configure the IP interface for your MTU size, 4344 bytes in my example:

```
ifconfig atm0 ipaddr netmask netmask mtu 4344
```

- Create a permanent ATM ARP cache entry for the ARP server with SDU size of 9188:

```
atmarp -s arpserver arpsrvnsap qos \
ubr:sdu=9188 arpsrv
```

- Configure the SDU (MTU plus 8 bytes) on the ATM interface:

```
atmarp -q network ubr:sdu=4352
```

Refer to Listing 1 for a complete example of configuring the ATM software for the Efficient adapter.

The configuration steps given are specific to IP over ATM connections using the Classical IP standard.

Using IP over Permanent Virtual Circuits

To establish a PVC, the following steps must be performed.

- On the workstation, add an ATM ARP entry on each node specifying the PVC (*vpi.vci* pair) used to connect to each of the other hosts.
- Create the PVC on the switch.

As an example, the following commands executed on the appropriate host will set up a PVC between nodes named node1 and node2, on interface 0, using a vpi of 0 and a vci of 70:

- `node1: atmarp -s node2 0.0.70`
- `node2: atmarp -s node1 0.0.70`

The PVC is identified by three numbers, separated by two periods. The numbering scheme is *interface.vpi.vci*, where interface is 0 for the first ATM adapter, 1 for the next, etc. The default interface for the atmarp command is 0. The *vpi* (virtual path identifier) and *vci* (virtual channel identifier) are the standard ATM PVC identifiers. The host name (node1 and node2) can be used if there is an entry for it in the `/etc/hosts` file; otherwise, use the IP address of the host.

The commands above tell node1 to communicate with node2 over PVC 0.0.70 and for node2 to communicate with node1 over PVC 0.0.70. The atmarp command links the IP address of the target host to the PVC. You could choose a different PVC for each connection, but it is simpler to think in terms of one PVC connecting two machines.

The *vpi.vci* pair must not be in use on the host. Also, any ATM ARP cache entries must be deleted for the target host before creating the PVC. (These cache entries are created when SVCs are opened to the destination host.) To delete an ARP cache entry on node1 for node2, you would use this command:

```
node1: atmarp -d node2
```

Next, the switch must be configured to complete the PVC between the hosts. It is helpful to understand the port naming convention used by the Fore switch. The port names consist of three identifiers:

- Board: the number of the switch board (same as the SCP number); each SCP controls one switch board.
- Network Module: the slot (A, B, C, or D) in the switch board containing the port.
- Port: the physical port number on the network module.

For example, port 1b3 refers to the first switch board, the second network module (module b) on board 1 and the third physical port on the second network module. The Fore ASX-200 switch has only one switch board, while the ASX-1000 switch has four. There is a maximum of four network modules per switch board and a maximum of six physical ports per module.

You must now create the virtual channels on the ATM switch. In our example, you would enter these commands on SCP 1:

```
localhost::> conf vcc
localhost::configuration vcc> new 1a1 0 70 1a2 0 70
localhost::configuration vcc> new 1a2 0 70 1a1 0 70
```

1a1 is the switch port for node1 and 1a2 is the switch port for node2.

The switch completes the PVC based on the input port to output port virtual channel connection (VCC) mapping. Note that the PVC *vpi.vci* (0.70) matches the *vpi.vci* given to the atmarp commands on the hosts.

The above commands will connect two ports on the same ATM switch board. The Fore ASX-1000 switch has up to four switch boards. If you are connecting machines on different switch boards, the procedure is more complicated, as you must connect each port to the switch fabric and connect the fabric to each port. Thus, if you wish to connect a machine on port 1a1 to a machine on port 3a1, the following commands are required:

On SCP 1:

```
localhost::> conf vcc
localhost::configuration vcc> new 1a1 0 70 1e3 0 70
localhost::configuration vcc> new 1e3 0 70 1a1 0 70
```

On SCP 3:

```
localhost::> conf vcc
localhost::configuration vcc> new 3a1 0 70 3e1 0 70
localhost::configuration vcc> new 3e1 0 70 3a1 0 70
```

On the Fore switch, the fabric connections are slot e. Therefore, port 1e3 refers to a connection from switch board 1 to switch board 3. Likewise, 3e1 refers to a connection from switch board 3 to switch board 1. Fore refers to these ports as “intra-fabric” ports.

Testing the Connections

Once the Classical IP setup is complete, all of the standard network tests can be performed. The simplest test is done by using the ping command to test the connection. One difference between SVC and PVC connections is a large latency for the first ping response when using SVCs. The reason for the latency is the setup time needed to establish the SVC. After the SVC is established, the latency for SVC and PVC connections should be the same.

After verifying the basic connectivity, you can run some network performance tests over the ATM connection. I have used the Netperf tool (see Resources) as well as some benchmarks developed locally. The maximum throughput performance is very good, around 132Mbps. This number is close to the maximum payload data rate for an OC-3 ATM network.

Conclusion

I have given instructions needed to set up the switch and hosts on an ATM network with Linux. The configuration steps given are specific to IP over ATM connections using the Classical IP standard. In addition to Classical IP, LAN

Emulation (LANE) can be used to carry IP over ATM. LANE is supported by the Linux-ATM software as well, but configuration of LANE is beyond the scope of this article. For more information, refer to the documentation in the Linux-ATM distribution.

Hosts can communicate in several other ways using an ATM interface without relying on Classical IP. The ATM software supports "native" ATM sockets, where applications can communicate directly over an ATM connection, bypassing the IP software completely.

If you are interested in learning about ATM technology but don't have ATM hardware, the Linux-ATM software can be of help. The software has

the capability to emulate an ATM device using TCP/IP to make the actual connection. By taking advantage of this support, you can get a head start on configuring ATM for Linux and learning the ATM programming interface.

Wayne J. Salamon is a Computer Scientist in the High Performance Systems and Services Division at the National Institute of Standards and Technology in Gaithersburg, MD. He has worked on system software for PCs, UNIX workstations and IBM mainframes for the past 12 years. When not doing computer stuff, he appears to play guitar, though only when connected to vacuum tube amplifiers. Wayne can be reached at wsalamon@nist.gov.

National Institute of Standards and Technology

Contributions by the National Institute of Standards and Technology (NIST) are not subject to copyright. Certain commercial products are identified here in order to document our experiments. Identification of such products does not imply endorsement by NIST.

Resources

"ATM on Linux, User's Guide", W. Almsberger, available at <http://lrcwww.epfl.ch/linux-atm/>.


"Netperf: A Network Performance Benchmark (Revision 2.0)", available at <http://www.cup.hp.com/netperf/NetperfPage.html>.

The device driver for the Fore PCA-200E ATM adapter is available at <ftp://os.inf.tu-dresden.de/pub/pca200e/>.

ForeRunner ATM Switch Configuration Manual, FORE Systems, Inc., March 1996.

ATM Theory and Application, McDysan, David E. and Darren L. Spohn, McGraw-Hill, 1995.

NEW!



MULTI-PORT
COM CARDS
ALSO AVAILABLE

Cool Card Tricks

FOR UNIX, NT AND CITRIX


PCI-RAS4

- PCI Plug & play compliance
- Four, integrated V.90-specification modems for data, fax and voice functions
- Drivers supplied for all major operating systems
- Lifetime Warranty


Getting Wired: A simple multimodem card that will enhance your connectivity.
The PCI-RAS4 is the easiest way to connect your server—or your entire network—directly to today's wired world. This card is the plug-&-play solution for your data, fax and voice messaging requirements. Remote users can interact online by remote login, utilizing file transfer, or using a remote network connection. The PCI-RAS4 integrates four internal 56K modems with a dedicated communications accelerator to give you the highest throughput while placing minimum load on your server's processor.

Its standards-based design works seamlessly with the widest range of operating systems (WinNT, Win95, SCO Unix...) and third party software. It's even managed by your standard OS tools and wizards—no learning curve. You're up-and-running in minutes.

Peace of Mind comes FREE with every Chase product.
The PCI-RAS4 becomes an integrated part of your server's security and communications system. You can leave your office every knowing that your firewall is secure and risk-free. Best of all, your card comes with Chase's Lifetime Warranty.



1-800-CHASE-US



1-800-237-8931

P.O. BOX 9069 • CLEARWATER, FL 34620-9839

For more information on Chase Research contact us at www.chaser.com, 800-242-7387 or info@chaser.com. Order the PCI-RAS4 from Tech Data's toll-free number.

w w w . c h a s e r . c o m

SAVE up to \$81

YES, Please start my subscription.

☐ **2 years (24 issues) for \$39** ☐ **1 year (12 issues) for \$22**
Regular \$120 newsstand cost Regular \$60 newsstand cost

Name _____

Title _____

Company _____

Address _____

City _____

State/ZIP/Country _____

Phone _____ Fax _____

E-mail _____

☐ Payment Enclosed ☐ Bill me Bill my card: ☐ VISA ☐ M/C ☐ AmEx

Card # _____

Exp. Date _____ Signature _____

Canadian and international subscriptions must be prepaid in U.S. funds

Please circle your foreign subscription choice:

Canada/Mexico rates: 2yr/\$49 1yr/\$27

International rates: 2yr/\$64 1yr/\$37

subs@ssc.com • www.linuxjournal.com • phone: 281.261.2581 • fax: 281.261.5999

Toll-free in the U.S. & Canada: 888-66-LINUX (888-665-4689)

PLEASE ALLOW 6-8 WEEKS FOR PROCESSING

58

LINUX JOURNAL

READER SERVICE February 1999

58

Expires 07/99

To request information from an advertiser:

Circle the corresponding RS #, and put this card in the mail or fax it to +1.281.261.5999

| | | | | | | | | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 |
| 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 | 51 | 52 | 53 | 54 |
| 55 | 56 | 57 | 58 | 59 | 60 | 61 | 62 | 63 | 64 | 65 | 66 | 67 | 68 | 69 |
| 70 | 71 | 72 | 73 | 74 | 75 | 76 | 77 | 78 | 79 | 80 | 81 | 82 | 83 | 84 |
| 85 | 86 | 87 | 88 | 89 | 90 | 91 | 92 | 93 | 94 | 95 | 96 | 97 | 98 | 99 |
| 100 | 101 | 102 | 103 | 104 | 105 | 106 | 107 | 108 | 109 | 110 | 111 | 112 | 113 | 114 |
| 115 | 116 | 117 | 118 | 119 | 120 | 121 | 122 | 123 | 124 | 125 | 126 | 127 | 128 | 129 |
| 130 | 131 | 132 | 133 | 134 | 135 | 136 | 137 | 138 | 139 | 140 | 141 | 142 | 143 | 144 |
| 145 | 146 | 147 | 148 | 149 | 150 | 151 | 152 | 153 | 154 | 155 | 156 | 157 | 158 | 159 |
| 160 | 161 | 162 | 163 | 164 | 165 | 166 | 167 | 168 | 169 | 170 | 171 | 172 | 173 | 174 |
| 175 | 176 | 177 | 178 | 179 | 180 | 181 | 182 | 183 | 184 | 185 | 186 | 187 | 188 | 189 |
| 190 | 191 | 192 | 193 | 194 | 195 | 196 | 197 | 198 | 199 | 200 | 201 | 202 | 203 | 204 |
| 205 | 206 | 207 | 208 | 209 | 210 | 211 | 212 | 213 | 214 | 215 | 216 | 217 | 218 | 219 |
| 220 | 221 | 222 | 223 | 224 | 225 | 226 | 227 | 228 | 229 | 230 | 231 | 232 | 233 | 234 |
| 235 | 236 | 237 | 238 | 239 | 240 | 241 | 242 | 243 | 244 | 245 | 246 | 247 | 248 | 249 |
| 250 | 251 | 252 | 253 | 254 | 255 | 256 | 257 | 258 | 259 | 260 | 261 | 262 | 263 | 264 |
| 265 | 266 | 267 | 268 | 269 | 270 | 271 | 272 | 273 | 274 | 275 | 276 | 277 | 278 | 279 |
| 280 | 281 | 282 | 283 | 284 | 285 | 286 | 287 | 288 | 289 | 290 | 291 | 292 | 293 | 294 |
| 295 | 296 | 297 | 298 | 299 | 300 | 301 | 302 | 303 | 304 | 305 | | | | |

Name _____

Title _____

Company _____

Address _____

City _____

State/ZIP/Country _____

Telephone _____ Fax _____

E-mail _____

Linux Journal has teamed up with S.u.S.E. to offer LJ readers a special promotion—subscribe (or renew your subscription) to *Linux Journal* today and receive the S.u.S.E. Linux 5.2 Sampler on CD-ROM FREE!*



Offering you German-engineered excellence, S.u.S.E. Linux 5.2 Sampler is a complete, easy-to-install Linux operating system which includes the KDE desktop and Netscape Communicator™. Get your free copy today while supplies last!

☐ **Yes!** Please send me a free CD with my pre-paid order.

*Offer available with pre-paid orders within the U.S.A. International orders: please add \$3.75USD for postage.

FREE S.U.S.E. Linux CD!

SAVE up to \$81

YES, Please start my subscription.

☐ **2 years (24 issues) for \$39** ☐ **1 year (12 issues) for \$22**
Regular \$120 newsstand cost Regular \$60 newsstand cost

Name _____

Title _____

Company _____

Address _____

City _____

State/ZIP/Country _____

Phone _____ Fax _____

E-mail _____

☐ Payment Enclosed ☐ Bill me Bill my card: ☐ VISA ☐ M/C ☐ AmEx

Card # _____

Exp. Date _____ Signature _____

Canadian and international subscriptions must be prepaid in U.S. funds

Please circle your foreign subscription choice:

Canada/Mexico rates: 2yr/\$49 1yr/\$27

International rates: 2yr/\$64 1yr/\$37

subs@ssc.com • www.linuxjournal.com • phone: 281.261.2581 • fax: 281.261.5999

Toll-free in the U.S. & Canada: 888-66-LINUX (888-665-4689)

PLEASE ALLOW 6-8 WEEKS FOR PROCESSING

58

TAPE EDGES, DO NOT STAPLE OR GLUE
Outside U.S. please place in envelope


BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 10, MISSOURI CITY, TX

POSTAGE WILL BE PAID BY ADDRESSEE

LINUX JOURNAL

PO BOX 500
MISSOURI CITY, TX 77459-9903
USA



LINUX JOURNAL

The Premier Linux Magazine

"It (Linux) is an invaluable learning tool, and *Linux Journal* is absolutely the best magazine I have found in the UNIX community. I can't tell you how many articles I have found applicable to not only Linux, but to other systems."

--- Matthew Huff

Linux Journal subscriber


"...LJ articles also introduce me to new things I would never have the time to research on my own..."

--- Richard Nolde

System Manager


<http://www.linuxjournal.com/>

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



LINUX JOURNAL

PO BOX 55549
SEATTLE, WA 98155-0549
USA



LINUX JOURNAL

PO BOX 500
MISSOURI CITY, TX 77459-0500
USA



PLACE
STAMP
HERE

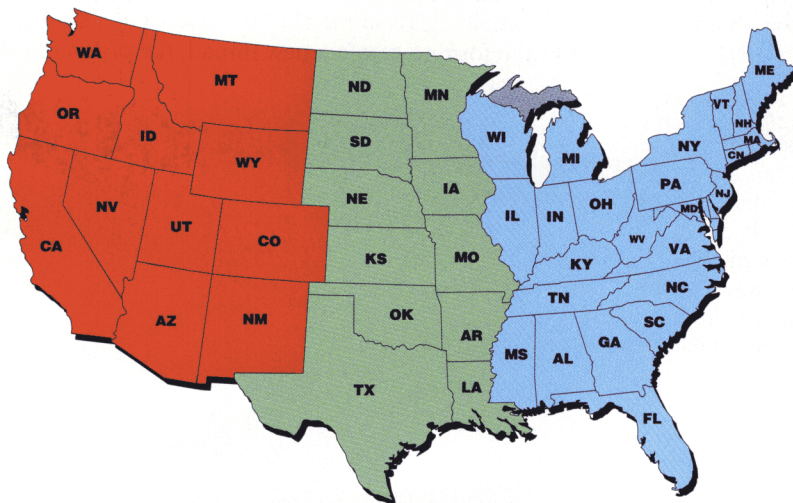
PLACE
STAMP
HERE

ADVERTISERS INDEX

| RS# | ADVERTISER | PAGE | RS# | ADVERTISER | PAGE |
|-----|--------------------------------------|--------|-----|-------------------------------------|----------|
| 261 | Appgen Business Software | 15 | 219 | Knox Software | 25 |
| 245 | ASA Computers, Inc. | 39 | 278 | Leverage Information Systems, Inc. | 58 |
| 184 | ASL Workstations, Inc. | 83 | 266 | Linux Central | 68 |
| 177 | Aspen Systems, Inc. | 5 | 243 | Linux Expo c/o Red Hat Software | 51 |
| 234 | Caldera, Inc. | 33 | 166 | Linux Hardware Solutions | 59 |
| 283 | Chase Research | 80 | 265 | Linux Journal 1997 Archive CD-ROM | 49 |
| 277 | Computone Corp. | 57 | 14 | Metro Link, Inc. | C4 |
| 265 | Computer Literacy Bookshops | 86 | 195 | Microway, Inc. | 7 |
| 284 | Consensys Computers, Inc. | 17 | 83 | Numerical Algorithms Group, Inc. | 71 |
| 33 | Cyclades Corporation | 11, 13 | 72 | O'Reilly and Associates | 27 |
| 218 | Cygnus Solutions | 31 | 207 | Perforce Software, Inc. | 35 |
| 128 | DCG Computers, Inc. | 91 | 279 | Professional Event Management, Inc. | 65 |
| 144 | Empress Software, Inc. | 23 | 210 | Raima Corporation | 67 |
| 55 | Enhanced Software Technologies, Inc. | 2 | 22 | Red Hat Software | 53 |
| 140 | FairCom Corporation | 29 | 24 | S.u.S.E. | 45 |
| 276 | Future Technologies | 89 | 269 | Softway Systems | 73 |
| 246 | Gcom, Inc. | 19 | 25 | SW Technology | 55 |
| 242 | Hi-Tech USA | 96 | 173 | Telenet System Solutions, Inc. | 61 |
| 255 | HKS, Inc. | 92 | 227 | The Portland Group | 69 |
| 264 | Hurricane Electric | 44 | 194 | Unitrends | 30 |
| 275 | IDG Linux World Expo | 37 | 12 | VA Research, Inc. | C2, 1, 9 |
| 270 | Image Stream Internet Solutions | 87 | 262 | VariCAD | 28 |
| 74 | InfoMagic, Inc. | C3 | 248 | Vital, Inc. | 75 |
| 280 | Inlimine, Inc. | 21 | 76 | Walnut Creek CDROM | 36 |
| 260 | InterBase Software Corporation | 47 | 188 | ZD Comdex & Forums | 85 |

LINUX JOURNAL

PO Box 55549 • Seattle, WA 98155-0549
Phone: 206.782.7733 • FAX: 206.782.7191
URL: <http://www.linuxjournal.com>



WESTERN USA

David Frankel
FRANKEL MEDIA SERVICES
Phone: 203.775.3928
E-Mail: dfrankel@ssc.com

MIDWESTERN USA

Carlie Fairchild
Phone: 206.782.7733
FAX: 206.782.7191
E-Mail: carlie@ssc.com

EASTERN USA

Emily Salus
Phone: 703.281.1013
FAX: 703.281.1013
E-Mail: emily@ssc.com

GERMANY/NETHERLANDS

Albert Ticheler
MEDIALTIC
Phone: +31.40.292.66.93
FAX: +31.40.292.66.98
E-Mail: altich@iaehv.nl

ADVERTISING AND MARKETING SERVICES

Carlie Fairchild
SALES AND MARKETING MANAGER
Phone: 206.782.7733
FAX: 206.782.7191
E-Mail: carlie@ssc.com

Anne Tiemann
ADVERTISING PRODUCTION COORDINATOR
Phone: 206.782.7733
FAX: 206.782.7191
E-Mail: ads@ssc.com

Rubin Response Service, Inc.
LIST MANAGEMENT
Phone: 847.619.9800 ext/139
FAX: 847.619.0149
E-Mail: kristenk@rubinresponse.com

Creating a Web-based BBS, Part 2



Mr. Lerner continues to look at the bulletin board system, examining the code that works with individual messages.

by Reuven M. Lerner

Last month, I demonstrated how to build a small bulletin-board system (BBS) on the Web using Perl and a relational database. Such a bulletin board is another useful tool for bringing people together on a web site. This month, I will show you how to write several different programs, including ones that create and list the current messages.

Before continuing, let's review the two tables that contain the information in our BBS. I used SQL to define the tables, which means that while I wrote the database using MySQL, most of these definitions should work with other relational databases as well. The code to create these two tables is shown in Listing 1.

These two tables will enable a number of tricks to be performed with threads (i.e., message groups) and messages. Each message belongs to one thread. Each message (and thread) is contained in a single database row that includes the author's name, her e-mail address, a subject heading and the text of the message.

Pointers to additional information about relational databases in general or MySQL in particular can be found in "Resources", which also includes information about Perl's vendor-independent database interface (DBI).

Creating a Thread with Cookies

Last month I wrote a program, `list-threads.pl`, to list the threads (discussion topics) currently defined, and one to create a new thread, `add-thread.pl`. However, I didn't provide an HTML form for use with `add-thread.pl`, because that form must be produced within a CGI program. The program that performs this function is `add-thread-form.pl` in Listing 2 in the archive file. Listings 2 through 5 are not printed here due to space considerations, but are available by anonymous download in the file <ftp://ftp.ssc.com/pub/lj/listings/issue58/3252.tgz>.

Why use a CGI program rather than a static form? To be honest, the only reason is to provide a bit of functionality I particularly like. I order many items on the Web, often returning to the same vendor multiple times. I dislike having to enter my name and e-mail address every time I fill out an HTML form on the Web. I decided to make life a bit easier for people using our bulletin board system by automatically filling in the "name" and "email" fields with information the user had posted in a previous transaction.

If I were to use a templating system such as Embperl or ePerl, I could use a file that looks closer to standard HTML

without burying the HTML inside of a CGI program. For a variety of reasons, including the fact that my web-space provider had not made `mod_perl` available as of this writing, I decided to use CGI programs rather than templates.

Regardless of whether CGI programs or templates are used, inserting a value from a previous form submission requires keeping track of state across HTTP transactions. HTTP is a stateless protocol; that is, each connection occurs without any memory from previous ones. How, then, can data be retrieved from a previous form submission?

The answer is HTTP cookies, a clever hack that has become a cornerstone of commerce and transaction on the Web. A cookie is a name,value pair, somewhat like a variable or an entry in a hash table. The cookie is stored by the user's browser, however, meaning that it is available across multiple transactions.

A site can set a cookie as part of an HTTP response, with a "Set-cookie" header. Whenever the user visits a site that previously set a cookie, it includes a "Cookie" header in its HTTP request. Thus, the cookie's value can be used to automatically fill in the "value" attribute of the "name" and "email" fields in the HTML form. When the user submits the form to create a new thread, the program sends

Listing 1. Table Creation Code

```
CREATE TABLE ATFThreads (
  id SMALLINT UNSIGNED AUTO_INCREMENT PRIMARY KEY,
  subject VARCHAR(255) NOT NULL,
  author VARCHAR(60) NOT NULL,
  email VARCHAR(60) NOT NULL,
  text TEXT NOT NULL,
  date DATETIME NOT NULL,
  UNIQUE(subject)
);

CREATE TABLE ATFMessages (
  id MEDIUMINT UNSIGNED AUTO_INCREMENT PRIMARY KEY,
  thread SMALLINT UNSIGNED NOT NULL,
  subject VARCHAR(60) NOT NULL DEFAULT "No subject",
  date DATETIME NOT NULL,
  author VARCHAR(60) NOT NULL DEFAULT "Mr. Nobody",
  email VARCHAR(60) NOT NULL DEFAULT
    "atf@lerner.co.il",
  text TEXT NOT NULL );
```


100

system bus

3-D AGP
graphics

up to

633

nhz

red hatTM
Linux

X_i GraphicsTM
Accelerated-X

ASL Workstations, Inc.

ASL Workstations, Inc. is the leading provider of high-end Linux workstations. With up-to-the-minute technology, our workstations offer superior performance and reliability. Comprised with highly-trained technical staff, ASL will innovate and custom-build the best workstation to fit your specification. Our dedication for performance and quality will blow your mind—not your budget.



For the latest systems configuration and pricing, please visit our web site at:

<http://www.aslab.com>

Voice: (408) 923 - 3837

Fax: (408) 923 - 4615

Email: sales@aslab.com



headers that set the "name" and "email" cookies on the user's browser. The next time the user visits the site, those values are sent as part of the HTTP headers and can be retrieved and used within our program.

Perl's CGI.pm module allows us to easily work with cookies, using the "cookie" method. The following code is put in the form-creation program:

```
my $email = $query->cookie(-name => "email") ||
    "";
print "<TR>\n";
print "<TD>Your e-mail address</TD>\n";
print "<TD><input type=\"text\" size=\"50\" ";
print "value=\"$email\" name=\"email\"></TD>\n";
print "</TR>\n\n";
```

which assigns \$email the value of the "email" cookie or the empty string. The empty string could be ignored, since Perl automatically assigns the empty string to a variable the first time its value is retrieved. However, this would produce a warning message, since the program would be using the value of an undefined variable. It is safest to assign the empty string when possible.

Each message (and thread) is contained in a single database row that includes the author's name, her e-mail address, a subject heading and the text of the message.

The value of the text field is set to the current value of the cookie or the empty string; that is, either the user's e-mail address or nothing at all. A similar method is used for the user's name, so that she doesn't have to enter her name multiple times.

The form is submitted to add-thread.pl, which we examined last month. That program uses the elements of the submitted HTML form to create an SQL query that inserts an appropriate row into the ATFThreads table. Because we have defined the ID column of ATFThreads with the AUTO_INCREMENT attribute, we can be sure every thread will have its own automatically generated ID number that we can reference in our programs.

When the form is submitted, our CGI program creates two new cookies, one named "email" and another named "name". We can then retrieve the values with CGI.pm's "cookie" method, as demonstrated above. Creating the cookies is almost as easy as retrieving them:

```
my $namecookie = $query->cookie(-name => "name",
    -value => $query->param("name"),
    -expires => "+1y");
my $emailcookie = $query->cookie(-name => "email",
    -value => $query->param("email"),
    -expires => "+1y");
```

Once we create \$namecookie and \$emailcookie, we can send them to the user's browser, thus setting the cookie values, by incorporating them into the HTTP header:

```
print $query->header(-type => "text/html",
    -cookie => [$namecookie, $emailcookie]);
```

Since both cookies are set to expire one year (+1y) after they are created, the user's browser should continue to send the name and e-mail address whenever visiting the site in the future.

Working with Messages

Now that we have seen how the underlying database system will work for threads, we need to begin working with the actual messages. Because this is a simple system, we'll look only at posting a new message to a thread and viewing the contents of a thread.

In many ways, posting a new message to a thread is similar to creating a new thread. In both cases, the user's name and e-mail address are requested. In both cases, the date and time at which the thread was created is recorded, and the user can enter a title and a message body.

The major difference between messages and threads is that each message must be associated with a thread. This association is used to create the illusion that the messages are stored separately, when in fact they are all stored in the same table (ATFMessages). Users, however, will be able to see only a single "thread-wise slice" at a time.

Just as I used a program to create the thread-adding form, I will use a CGI program to create the message-posting form, called post-comment-form.pl (Listing 3 in the archive file). This form will submit its contents to post-comment.pl (Listing 4 in the archive file).

I will ensure that each message is associated with a thread by putting a selection list inside of the form. Each option in the selection list will be identified internally with the ID code for the thread in question and will display the subject line.

In order for this list to reflect the current status of the database, a database query is done and the results are displayed in the form. The query is set by:

```
my $sql =
    "SELECT id,subject FROM ATFThreads ORDER BY subject";
```

and then executed, iterating through each id, subject pair. Each pair is inserted into an <option> tag, as we can see:

```
while (my @row = $sth->fetchrow)
{
    print "<option value=\"$row[0]\" ";
    print " selected " if ($thread_id == $row[0]);
    print ">$row[1]\n";
}
```

The standard DBI \$sth->fetchrow method is used to return the next row from the SELECT query. When no more rows remain to be retrieved, \$sth->fetchrow returns false, which ends the while loop.

Also notice how a particular thread's subject can be selected by comparing its \$thread_id with \$row[0]. \$thread_id is set to the value of the query string, which can be loosely defined as "anything following the question mark in a URL". The line:

```
my $thread_id = $query->param("keywords") || 0;
```

causes CGI.pm to automatically assign the parameter keywords to the value of the query string. If the user invokes the program with <http://www.lerner.co.il/cgi-bin/post-comment-form.pl?5>, then \$thread_id will be assigned the value 5. If the query string is not assigned, the value is left at 0, in which case no default thread is selected.


Next Stop, Chicago.

*The season may change,
but the IT issues change even faster.*

DIGITAL CONSUMER TECHNOLOGY



elecom
remote access



Internet
co
Starbucks
General Motors
For Chrysler
AMOC
Alstria
Sara Li
UA
Northwestern Mutual
Life Insurance
Walgreen
Deer
Abbott Laboratories
WMX Technologies
First Chicago NBD Corp
Farmland Industries
Johnson Controls
Unicom
R.R. Donnelley & Son
Manpower
Navistar International
Baxter International
Cass
Quaker Oats
Stone Container
Illinois Tool Works

COMDEX
Spring '99

**Windows
World**

Become a part of the
Linux Global Summit
at COMDEX/Spring '99.

Sponsored by Linux International

April 19-22, 1999 • McCormick Place • Chicago, Illinois

For information on attending,
call the Conference Hotline at 781-433-1650
or go online at www.comdex.com.

Also at COMDEX/Spring and WINDOWS WORLD is the Linux
Pavilion. Companies interested in exhibiting should call Bill
Mahan at 781-433-1817 or e-mail sonny_saslaw@zd.com.

Posting the Message

When the HTML form is submitted to `post-message.pl` (Listing 4 in the archive file), the form elements are used to insert a new row into `ATFMessages`. As I indicated above, `post-message.pl` is not very different from `add-thread.pl`, except that it stores a thread ID number along with all the other information:

```
my $sql = "INSERT INTO ATFMessages ";
$sql .= " (thread,date,author,email,subject,text) ";
$sql .= "VALUES ";
$sql .= " ($thread_id,NOW(),$name,$email,$subject,$text)";
```

The variable values can be inserted without surrounding them by quotes, because the standard `$dbh->quote` method was used. I discovered this method only recently and continue to be amazed that I was ever able to survive without it. Simultaneously, the form elements are retrieved and quoted appropriately in the following lines of code:

```
my $name = $dbh->quote($query->param("name"));
my $email = $dbh->quote($query->param("email"));
my $thread_id = $dbh->quote($query->param("thread"));
my $subject = $dbh->quote($query->param("subject"));
my $text = $dbh->quote($query->param("text"));
```

Once this is done, the above SQL query will INSERT a new row. We tell the user that the new message has been added and produce a menu bar with a number of options.

Believe it or not, these two short programs are all that is needed to insert a message into the database and thus into our BBS.

Viewing a Thread

At this point, the functionality is close to complete. All that remains to be done is to create `view-thread.pl` (Listing 5 in the archive file), which allows us to look at the current contents of a thread.

For this program to work, a single argument must be passed in the query string to identify the thread. To retrieve this value, use the keywords HTML form element that CGI.pm creates:

```
my $thread_id = $query->param("keywords");
```

Once `$thread_id` is assigned, I can retrieve the appropriate information from the tables about that thread. Indeed, two separate queries are done: one from `ATFThreads` and a second from `ATFMessages`. (I could have combined the queries into a single large SELECT statement, but I chose to keep them separate.)

Early on, I decided to print the date and time of the user's posting along with the text of the posting. Given the `DATETIME` data type, how can we retrieve the date and time in an intelligent way? MySQL provides a `DATE_FORMAT` function which takes the value from a column and writes the contents using a specified format.

To make life easier, I actually retrieve the same "date" column twice, once for the date and again for the time. This allows literal characters to be inserted between the date and time without having to worry about possible misinterpretation:

```
$sql = "SELECT id, DATE_FORMAT(date,
    \"%W, %d %b %Y\"), ";
$sql .= "DATE_FORMAT(date, \"%h:%i %p\"), ";
$sql .= "author, email, subject, text FROM ATFMessages ";
$sql .= "WHERE thread = $thread_id ORDER BY date desc";
```

`DATE_FORMAT` takes two arguments: the name of the column to retrieve and a set of codes (in the style of C's `printf` statement) indicating the values to use.

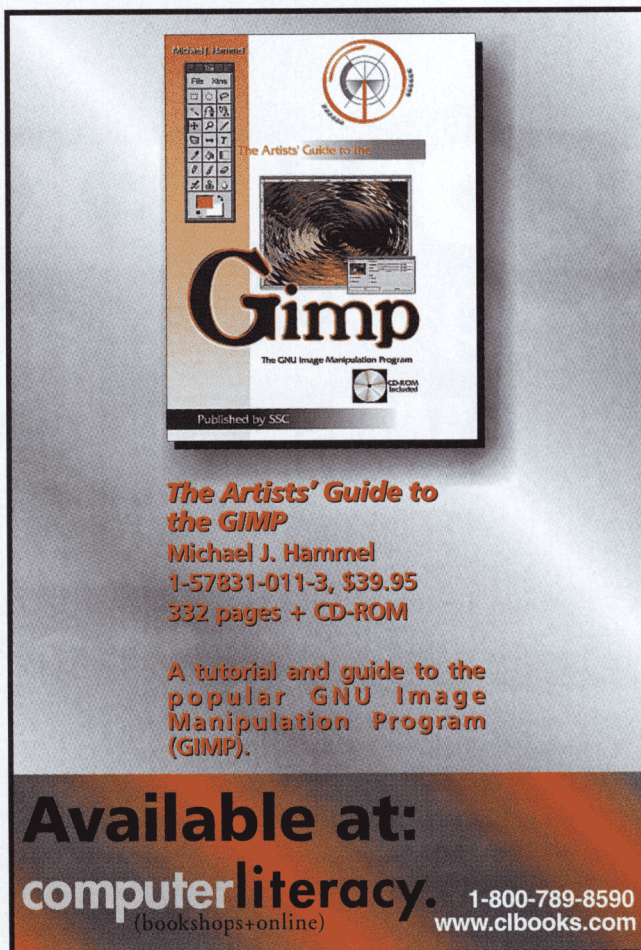
Once this query is executed, the code iterates through the results, printing the messages as they come—from newest to oldest. They will come in that order because of the `ORDER BY` clause in the SELECT statement. Allowing the database to do our dirty work for us means we can print all of the messages in a thread with just the following short loop:

```
while (my @row = $sth->fetchrow)
{
    my ($id, $date, $time, $author, $email, $subject,
        $text) = @row;
    print "<a name=\"$id\"><B>$subject</B>, ";
    print "by <a href=\"mailto:$email\">$author</a> ";
    print "on $date at $time</P>\n";
    print "<blockquote>$text</blockquote>\n\n";
}
```

Summary

The basic BBS software is now finished. It can create threads, add a message to a thread and view the messages within a thread. If nothing else, this project shows how powerful a set of database tables can be when paired with some CGI programs.

Perhaps this system is too basic; it is lacking some functions that any good bulletin board (or any good web application) should handle. For instance, it would be nice to include the ability to search through the posted messages for a text string or regular expression to find messages relevant to a particular topic.




The Artists' Guide to the GIMP
Michael J. Hammel
1-57831-011-3, \$39.95
332 pages + CD-ROM

A tutorial and guide to the popular GNU Image Manipulation Program (GIMP).

Available at:
computerliteracy. 1-800-789-8590
(bookshops+online) www.clbooks.com

It would also be useful to provide some administrative functions, unavailable to the public at large, for handling things at a relatively high level. For example, we might want to reserve the ability to delete messages that are offensive or unrelated to the topic. We might even want to give this ability to certain other users, giving them "deputy moderator" status.

Next month, we will see how to provide these functions by adding just one or two new CGI programs. Meanwhile, you can see these programs in action at <http://www.lerner.co.il/atf/>, where the "At the Forge" BBS is already in use. 



Reuven M. Lerner is an Internet and Web consultant living in Haifa, Israel, who has been using the Web since early 1993. In his spare time, he cooks, reads and volunteers with educational projects in his community. You can reach him at reuven@netvision.net.il.

Resources

Perl: The Perl home page is at <http://www.perl.com/> and includes links to distributions, tutorials, documentation, frequently asked questions and more. This is a good launching pad for information about Perl and for keeping up to date on the latest versions of Perl and its modules.


DBI: The Perl "Database Interface", a set of modules to provide a generic API (DBI) and drivers to specific databases (DBD), allows for portable access to a large number of databases. Information about DBI is available from the Perl home page, as well as from <http://www.hermetica.com/tecnologia/perl/DBI>.

Apache: The Apache web server, used on a majority of web sites, is freely available from <http://www.apache.org/> in source format. Documentation for Apache, as well as tutorials on how to use and administer it, can be found at that web site or <http://www.apacheweek.com/>, a weekly newsletter for Apache users.

MySQL: The MySQL home page is at <http://www.mysql.com/>, although there are numerous mirrors around the world. Some additional MySQL documentation can be found at <http://www.turbolift.com/>.

Database-Backed Web Sites, a book by Phil Greenspun on the creation and maintenance of databases and web sites, is a classic that should be read by anyone using these tools. See <http://photo.net/wtr/> for a full copy of the book (as well as a draft of the coming sequel), along with many other useful tidbits. While I did not use his code when designing this system, I did use some of his ideas about formatting and presentation. So while the systems don't share any code, they do have a similar "look and feel".

WE'RE PLAYING WITH A FULL DECK.



High-Performance, High-Reliability
WAN CARDS FOR LINUX
Comprehensive, Scalable Solutions from
ImageStream & SDL Communications

**56/64 Kbps ■ Fractional & Full T1/E1
Fractional & Full T3/E3 up to 52 Mbps**

PPP ■ HDLC ■ Frame Relay ■ Raw IP

155 Mb ATM over OC3 Coming Soon!

| ISA | PCI | CompactPCI |
|--------------|-----------|------------|
| RISCom/N2s | WANic 400 | Aries 600 |
| RISCom/N2d | WANic 405 | Aries 650 |
| RISCom/N2dds | WANic 600 | Aries 800 |
| RISCom/N2csu | WANic 650 | Aries 805 |
| | WANic 800 | |
| | WANic 805 | |

ISis ImageStream™
Internet Solutions

Phone: (219) 935-8484 sales@imagestream-is.com
Sales: (800) 813-5123 www.imagestream-is.com
Fax: (219) 935-8488 ftp.imagestream-is.com

Ask about our reseller program.

G.L.U.E. = Groups of Linux Users Everywhere:

<http://www.ssc.com/glue/>

iHTML Merchant 2.0

Inline Internet Systems, Inc. announced the availability of the iHTML Merchant 2.0 software solution for electronic commerce on the Internet. iHTML Merchant 2.0 enables business owners, web developers and ISPs to deploy sophisticated on-line storefronts. Upgrades of the software from 1.0 are \$149US. Version 2.0 is \$739US per web server, which includes iHTML Pro 2.16 (available separately at \$590US).

Contact: Inline Internet Systems, Inc., 7305 Rapistan Court, Mississauga, ON L5N 5Z4, Canada, Phone: 905-813-8800, Fax: 905-813-8286, E-mail: info@inline.net, URL: <http://www.ihtmlmerchant.com/>.

Auto-Changer

UniTrends Software has released Auto-Changer, tape library software for Linux. Auto-Changer allows Linux users to back up and work with tape libraries. CTAR sells for \$195 on Linux (Intel) and Auto-Changer is \$149 for Linux, but must be used with CTAR, the Business Critical Backup Solution.

Contact: UniTrends Software Corporation, 1601 Oak St., Suite 201, Myrtle Beach, SC 29577, Phone: 800-648-2827 (or 843-626-2878 outside the US), Fax: 843-626-5202, E-mail: sales@unitrends.com, URL: <http://www.unitrends.com>.

JetStor RAID

AC&NC announced the JetStor RAID desktop-sized SCSI disk array system which includes a new high-performance RAID controller as well as functional improvements. Complete information including technical specifications can be found on Advanced Computer and Network Corporation's web site. Sample pricing: 32GB for \$6,500, 63GB for \$7,999 and 126GB for \$12,250.

Contact: AC&NC, 5001 Baum Blvd., Suite 770, Pittsburgh, PA 15213, Phone: 412-683-9010, Fax: 412-683-9070, E-mail info@acnc.com, URL: http://www.acnc.com/product_jetstor.html.

VariCAD 7.0-1.0

VariCAD has released a new version of its professional CAD system for mechanical engineering—VariCAD 7.0-1.0. VariCAD 7.0-1.0 is able to communicate with rendering software and FEM software (Cosmos, NuGraf, etc.). Free trial and demo versions are available at VariCAD's web site. The prices of the VariCAD system remain unchanged—VariCAD for Linux is \$299-\$499.

Contact: VariCAD, 931 Greenbriar Avenue, Ottawa, ON K2C 0J8, Canada, Phone: 613-723-0953, E-mail: info@varicad.com, URL: <http://www.varicad.com/>.

MetaCard 2.2

MetaCard Corporation announced the release of MetaCard 2.2. New features include support for new image formats including PNG and animated GIFs, the ability to resize images on all platforms, and many new properties, commands and functions to make application development faster and easier. The free MetaCard 2.2 Starter Kit is available from the MetaCard web site.

Contact: MetaCard Corporation, 4710 Shoup Pl.,



Boulder, CO 80303, Phone: 303-447-3936, Fax: 303-499-9855, E-mail: info@metacard.com, URL: <http://www.metacard.com/>.

Macsyma 421 for Linux

Macsyma Inc. announced the release of Macsyma 421 math software for Linux. Macsyma offers mathematical power in symbolic, numerical and graphical mathematics. The PC interface offers Macsyma's MathTips Advisor, which allows users to type questions in their own words and receive executable "tips". The U.S. commercial price for Macsyma 421 for Linux workstations is \$249 (or \$199 without paper manuals).

Contact: Macsyma Inc., 20 Academy Street, Arlington, MA 02476, Phone: 781-646-4550, Fax: 781-646-3161, E-mail: info@macsyma.com, URL: <http://www.macsyma.com/>.

HP Firehunter Supports Linux

Hewlett-Packard Company announced that its HP Firehunter family of Internet service-management solutions, targeted for the Internet service provider (ISP) community, now supports Red Hat Linux. HP Firehunter is a family of measurement and monitoring solutions designed specifically to help ISPs proactively manage Internet services such as mail, news and web functions. The product family includes Firehunter/L (\$1,450) for small ISPs running up to 20 Internet servers and Firehunter 1.5 for mid-sized providers with up to 60 servers. See web site for pricing.

Contact: Hewlett-Packard Company, 3000 Hanover Street, Palo Alto, CA 94304, Phone: 650-857-1501, URL: <http://www.firehunter.com/>.

PlanetUplink

Planet Computer's newest business solution, PlanetUplink, has been expanded to support Linux (server and client). PlanetUplink IBN (Internet Based Network) allows businesses to gain access to and share virtually any application or database simultaneously (real time) on almost any computer from their remote and multiple offices, globally, via the Internet. IBN is available in 2 versions: Internet Office Hosting (IOH) and Internet Application Hosting (IAH). IAH/IOH user costs vary from \$45 to \$75 per month, in addition to setup fees depending upon the application(s).

Contact: Planet Computer, 910 16th St., Suite 624, Denver, CO 80202, Phone: 303-825-1778, Fax: 303-825-1773, E-mail: planet@planet-computer.com, URL: <http://planetuplink.com/>.

NetTracker 3.5 and Webmaster 3.0

Sane Solutions, LLC and COAST Software Inc. announced the offer of a package combining the NetTracker 3.5 Professional web-based usage tracking software and the COAST Webmaster 3.0 web site management software. The NetTracker 3.5 product is designed to provide a simplified solution for web site traffic analysis. COAST WebMaster 3.0 is used to monitor and maintain the content of Internet and Intranet web sites. The NetTracker 3.5 Professional/COAST WebMaster 3.0 bundle is available for \$750US.

Contact: Sane Solutions, LLC, 35 Belver Ave., Suite 230, North Kingstown, RI 02852, Phone: 800-407-3570, E-mail: info@sane.com, URL: <http://www.sane.com/>.

COAST Software Inc., 60 Queen Street, 14th Floor, Ottawa, ON K1P 5Y7, Canada, Phone: 613-567-3201, E-mail: info@coast.com, URL: <http://www.coast.com>.

ICE.PPN

J. River has released ICE.PPN (Portable Private Networking) to provide secure, end-to-end transmission of data between a Windows client and a corporate UNIX server, regardless of whether a firewall is in place. A full-featured 30-day trialware version of ICE.PPN can be downloaded from J.River's web site. Suggested list price for a server license with 5 clients is \$1195. Reseller pricing is available.

Contact: J. River, Inc., 125 North First Street, Minneapolis, MN 55401, Phone: 612-677-8200, E-mail: info@jriver.com, URL: <http://www.jriver.com/>.

O2 Object Database Management System

Ardent Software, Inc. announced the availability of the O2 Object Database Management System (ODBMS) for the Linux operating system. Ardent's O2 System, an ODMG-compliant ODBMS available on Linux, gives developers access to language bindings, OQL (Object Query Language) and O2Web, an integrated set of tools for simple and rapid development of web applications.

Contact: Ardent Software, Inc., 50 Washington Street, Westboro, MA 01581-1021, Phone: 508-366-3888, Fax: 508-366-3669, E-mail: info@ardentsoftware.com, URL: <http://www.ardentsoftware.com/>.

Quant-X and Altera SQL Server version 2.0

The Altera SQL Server is a complete transactional data storage and retrieval system based on the relational database model and also a programmable web server. It is a multi-user, relational database with ODBC and JDBC connectivity, demanding relatively few resources. Quant-X has provided Altera with the appropriate hardware in order to port the Altera SQL Server on the above-mentioned platform. Quant-X and Altera came to an agreement regarding the bundling of Altera SQL Server Version 2.0 with Quant-X hardware. See web site for pricing.

Contact: Altera Ltd., E-mail: webmaster@altera.gr, URL: <http://www.altera.gr/>.

Linux Edition of Ingres II

Computer Associates International, Inc. announced the activation of its Ingres II Linux Edition Open Beta program. This program enables customers to leverage the Linux platform in building core business applications in an *n*-tier environment, as well as preview the new version of CAI's RDBMS. The beta version of Ingres II Linux Edition can be downloaded for free from <http://www.cai.com/products/betas/>. When the generally

available version is released, it will also be offered as a free download.


Contact: Computer Associates International, Inc., One Computer Associates Plaza, Islandia, NY 11788, Phone: 1-888-7INGRES (888-746-4737), E-mail: info@cai.com, URL: <http://www.cai.com/>.

Thin Client PC X Server

GraphOn announced the world's first thin-client PC X server solution, providing high performance access to the X Window System and UNIX-based applications anywhere on an organization's Intranet, the public Internet or over dial-up. Visit the web site for pricing and further information.

Contact: GraphOn Corporation, 150 Harrison Avenue, Campbell, CA 95008, Phone: 408-370-4080, Fax: 408-370-5047, E-mail: sales@graphon.com, URL: <http://www.graphon.com/>.

Please send information about releases of Linux-related products to Ellen Dahl at newproducts@ssc.com or New Products c/o *Linux Journal*, P.O. Box 85867, Seattle, WA 98145-1876, or fax to 206-782-7191. Submissions are edited for length and content. Press releases are acceptable, but for fastest processing send a one-paragraph product description with contact information. Submissions can also be made on our web site (see <http://www.linuxresources.com/news/cool.html>) where they will also be posted.



Do you want to learn Linux?

Try
FTLinuxCourse!
Available Now!

FTLinuxCourse is a Course for Linux based on Caldera OpenLinux, written in HTML available on a CD-ROM, in your favourite language: English, Spanish, German, French and Italian.

It covers:

- OpenLinux Installation Step-by-Step.
- StarOffice, Communicator and KDE tutorials.
- 20 Chapters to introduce you to all Linux topics.
- Complete Linux Command References with examples On-Line.
- More than 500 questions with answers including 50 tests for Linux.
- All examples from the O'Reilly & Associates books.
- ...it also includes commercial demos from Xi Graphics, Helius, the KDE, all OpenLinux updates and more!

Browse the BASE course, at <http://www.futuretg.com/FTLinuxCourse/>
For orders & information contact us: sales@futuretg.com

LOOKING FOR AN EMPLOYEE WITH LINUX
EXPERIENCE? LOOKING FOR A JOB OPPORTUNITY
WHICH UTILIZES YOUR LINUX EXPERIENCE?
<http://www.linuxresources.com/employ/index.html>

Pre-installed Slackware

I picked up a used laptop that has Slackware partitioned on the hard drive. Is there anyway to log in to it without the password? I have a UNIX book but it does not address this problem. Do I have to re-install? I have no idea who the previous owner was.

--
Mike
shbecke@ibm.net

Try to boot the computer in single user mode (type `linux single` at the LILO prompt). If this does not work, you should get boot floppies in order to start a minimum LINUX in a RAM disk. Such floppies should be available in a Slackware distribution. Then you could mount the root partition by typing:

```
mount -t ext2 /dev/hda1 /mnt
and edit /mnt/etc/passwd in order to suppress the root password.
```

--
Pierre Fichoux, pfichoux@com1.fr

One of the simplest ways to do this is to obtain a boot and root disk pair from the Slackware distribution. Boot using both and use `mount(8)` to mount the root file system from your laptop. You can then edit the `mount_dir/etc/passwd` (or `mount_dir/etc/shadow`) file to remove the password.

For system administrators concerned with this potential security problem, the only way to prevent it is through the use of BIOS features that prevent access to the floppy drive without a password. This is not a Linux security flaw—any UNIX platform that may be booted from a set of floppies or from a CD-ROM that allows access to the hard drive has the potential to be “hacked” in this way. Notable exceptions are systems that use encrypted file systems, but those are rare and often much slower during normal operation.

--
Chad Robinson
chadr@brt.com

Software in RPMs

I have Slackware Linux. Wherever I look, I find software that is in RPMs. Is there no option other than shifting to Red Hat? Any conversion utilities?

--
Aseem Asthana
asthana@bom4.vsnl.net.in

Try the package <http://www.chez.com/imil/stuff/rpm4everyone.targz>; I use it every day on an old Slackware 3.0 distribution.

--
Pierre Fichoux
pfichoux@com1.fr

*Try **alien**, a package that converts packages from one format to another. Otherwise, compile **rpm** on your system so you can do the following:*

```
rpm2cpio package.rpm | cpio -list
rpm2cpio package.rpm | cpio -make-dirs -extract
```

It is quite handy. Note, however, that you might find incompatibilities between your system and the Red Hat packages. I always prefer compiling programs from source when I am on Slackware (but I might be overly cautious).

--
Alessandro Rubini
rubini@prosa.it

a.out

I just installed Red Hat's Linux 5.0. I used the C compiler to compile a simple C program, test.c, using the command:

```
gcc test.c
```

The program compiled and produced an executable called a.out. When I try to run a.out by typing:

```
a.out
```

I get a message that says a.out is not a recognized command. What am I doing wrong?

--
Jeff Miller
jeff_miller@msn.com

First confirm that the a.out file has the correct permissions. Use `ls -al a.out` to confirm that the executable (x) bit is set. If it isn't, use the `chmod +x` command to set this flag. If the permissions are correct, specify the full path to the file, as the current directory usually isn't in your default path. Use `./a.out` to ensure you are attempting to execute the correct program. Change your path by editing the `/etc/profile` file or the profile file in your home directory.

--
Vince Waldon
vince.waldon@cnpl.enbridge.com

Disk Space

How do I find out how much free space is left on my disk?

--
Kirk
sci@wadi-petro.com

*Use the **df** command. Briefly, running **df** without arguments will show the free space (in KB) on all your disks; `df /some_directory` will show the space left on the disk that `/some_directory` is on. Also note that **df** has an undocumented `-h` option that shows sizes in GB or MB as appropriate—convenient for today's large disks. See the man pages for information on other options.*

--
Scott Maxwell
s-max@pacbell.net

Relaying E-mail

The problem is the following. Each time I try to send an e-mail using a program that manages pop3 accounts such as Eudora or Netscape Mail, I receive the following message: "The recipient user@domain.com is not acceptable to your SMTP server. The message is not sendable until the recipient has been changed."

This problem appeared after we upgraded from the previous version of Linux to version 5.0. No problem occurs when receiving e-mail using these programs or when sending e-mail through Pine—only when using Eudora or any other similar program. How can we solve this?

--
Ricardo A. Williams L.
rick@corotu.stri.si.edu

The problem is in the new security policies of Red Hat 5. Your mail gets refused with a message of "551 we do not relay". The solution here is authorizing your client machine to relay mail through the Linux server. Your `/etc/sendmail.cf` is quite clear about the options:

```
# file containing IP numbers of machines which
# can use our relay
F{LocalIP} /etc/mail/ip_allow
```



```
# file containing names of machines which can
# use our relay
F{LocalNames} /etc/mail/name_allow
# file containing names we relay to
F{RelayTo} /etc/mail/relay_allow
```

Add lines to the proper file describing either the client's IP, the client's name or the recipient's name.

--

Alessandro Rubini
rubini@prosa.it

Digital UNIX Gateway?

I currently have a TCP/IP (via modem) connection between my Linux box at home and my office workstation (DEC station running Digital UNIX). The problem is that the DEC machine is *not* a gateway, so I cannot reach the rest of the subnet or the rest of the world, for that matter. Is there a way my Linux box can reach the subnet gateway which is two hops away? The **route** command in my current version of Linux (Slackware, kernel 2.0.0) does not support the **-hobcount** flag, which is supported by Digital UNIX and would do the trick.

--

Martin Olivera
martin@pantano.ucsd.edu

If your PPP IP address is one from the subnet your DEC station is sitting on, you just need to make sure it does ARP proxying for your Linux machine (in other words, it has to accept packets for your Linux machine's IP on its local Ethernet). If this is not the case, then it is more difficult. The options you have are:

- Find out if Digital UNIX can do IP masquerading like Linux can.

- Configure **routed** on your DEC server and advertise a route pointing to your Linux machine. Note that this will not work if your default gateway ignores RIP information, and it may upset your network administrator and/or be against company policies.

--

Marc Merlin
marc@merlins.org

Typing Spanish Characters

As a Spanish speaker, I want to use a keyboard with a complete set of Latin characters. I succeeded in implementing it for almost all applications, except Netscape. I installed the XKeysymDB file in the correct place, and this file works properly for Sun machines (I tested), but not for PCs with Linux. I tried to find the answer at Netscape's home page, but I couldn't. Perhaps I did the wrong search. Does anyone know how to set Netscape in order to have a "compose" key which produces accents, tildes and all that sort of thing?


--

Guigue
guigue@nucate.unicamp.br


If I'm not mistaken, the XKeysymDB file works only for a particular keyboard, so the one that works for your Sun keyboard is unlikely to work for your Linux machine's keyboard. Jamie Zawinski's **xkeycaps** found at <http://www.jwz.org/xkeycaps/> may help; it is a graphical editor for editing keyboard setups under X.


--

Scott Maxwell
s-max@pacbell.net



Full Support of Beowulf Linux





DCG Computers now offers Alpha, Intel & Sun Ultra Sparc



Support


| | | |
|---------|--------------|------------|
| Linux | Digital Unix | Windows NT |
| Solaris | Open VMS | |

Tel: 603-421-1800 Fax: 603-421-0911

For a complete list of Linux Solutions visit us at www.dcginc.com

All trademarks and registered trademarks are the sole property of their respective owners



Migrating from Windows

I got interested in Linux through a programme on BBC World. However, after an afternoon roaming the Internet, I couldn't get an answer to two questions most laymen probably have on the subject: does Linux replace Windows on my computer and is the process irreversible; will I be able to use my Windows-based programmes on Linux?


Philippe Humblé
humble@mbox1.ufsc.br

Linux can replace Windows on your computer, or the two can coexist. If you buy a commercial distribution or a Linux book, it should help explain how to do this. The process is irreversible only in the sense that you'll never want to go back.

There are different ways to use your Windows-based software under Linux. A couple of emulators—Wine and Wabi—enable you to run some Windows software directly under Linux. (Similarly, DOSEMU lets you run DOS software under Linux.) Other Windows software can be run by rebooting the machine into Windows, using the software, and then rebooting into Linux again as soon as possible.

As time goes on, you'll discover Linux software that can partially or entirely eliminate your need for Windows—Linux-native word processors, spreadsheets and such.

Scott Maxwell
s-max@pacbell.net



Alternatives to X

I am new to Linux and have heard much about the X Window System when word processing using Applixware Office or Corel's Word Perfect for Linux. What if I don't want to use X? Are there quality office applications that will run without X?

John Tam
johta@mailexcite.com

There is not, as far as I know, a non-X integrated office suite, but many of the pieces exist. For document processing, you can use TeX. Whether you'll like TeX or not depends on your needs—it is rock-solid and extremely powerful, but it is not WYSIWYG.

Scott Maxwell
s-max@pacbell.net

Most people associate "quality office applications" with "graphical user interface", a "what you see is what you get" environment like the current flock of office applications shipping for Windows machines. On UNIX systems, the standard GUI environment is the X Window System, so all of the current office suites for Linux run on top of X. This makes life much easier for the programmer—she can concentrate on writing a good word processor and leave the details of making it "graphical" to X. Early versions of X were somewhat tricky to set up and supported only a small subset of available graphics cards, but the configuration programs have come a long way as has driver support—a quick read of the X HOWTO should leave you with nothing to fear from "getting graphical" on your Linux machine.

Vince Waldon
vince.waldon@cnpl.enbridge.com

Move Money with Unix

Select a payment option:

- ☐ Credit Card
- ☐ Check
- ☐ Automatic Withdrawal



You've been accepting credit cards with Credit Card Verification System. Now you can accept checks with CCVS, or even take direct withdrawal from your customers' bank accounts. It's never been easier for your customers to give you money!

Take payments by check: CCVS provides **check verification and guarantee** to protect you from check fraud. Now, using ACH (Automated Clearing House), CCVS also gives you the

power of **direct withdrawal**, so you can move money directly from your customers' bank accounts without a paper check.

Clear credit cards with Unix: No hidden costs. No per-transaction fees. Free upgrades and free support for 6 months. Free developer's toolkit with APIs in C, Tcl, Perl, Java, Python, and PHP for easy integration.

CCVS: the ultimate E-Commerce solution.

Questions? Visit our Web site at
<http://www.hks.net/>
or email sales@hks.net
2732 Murray Avenue
Pittsburgh, PA 15217
1-888-HKS-7836 voice
1-412-521-2994 fax



Many of the on-line resources referred to in this column are available on the SSC web pages. Sunsite mirror sites, FAQs and HOWTOs can all be found at <http://www.linuxresources.com/>. Answers published in Best of Technical Support are provided by a team of Linux experts.

If you would like to submit a question for consideration for use in this column, please fill out the web form at <http://www.linuxjournal.com/techsup.html> or send e-mail with the subject line "BTS" to Ellen Dahl at bts@ssc.com.

Occasionally *Linux Journal* makes its mailing list available to vendors who have products or announcements relevant to the Linux community. If you prefer that your name not be included in vendor mailings, please let us know. Write to: *Linux Journal*, P.O. Box 500, Missouri City, TX 77459-9903 with the information from your mailing label.



The login Process

The beginning of it all...

by Andy Vaught

Virtually all Linux sessions begin with the user typing his user name at a prompt that looks like this:

login:

In this article, I will explain a little about what really happens behind the scenes and what contortions the system goes through to get a user going.

A Bit about the Shell

First, a quick look at the shell. The shell, which is just a program like any other, reads the characters you type and looks for a program with the same name. Program names are typed at the prompt and executed by the shell. Ending a command line with the & character causes the command to be run in the background.

The shell runs a program in two steps. First, the shell does an operation called a “fork”. Forking creates a new process that looks just like the original process, inheriting many attributes of its parent such as any open files and user ID. Although it is an exact copy of the shell program, the “child” process does not read user commands. The child shell immediately does an operation called an “exec”, short for “execute”, in which it causes the Linux kernel to load the new program over the top of the child shell and run that program in its place.

At this point, the original shell simply waits for the child program to finish. Once done, it gets the next line of input from the user, then the whole procedure is repeated. In an active UNIX system, this sort of thing is happening all the time. Even on fairly inactive systems, processes are still run to do housekeeping chores, while others are simply sleeping and waiting for something to happen.

From the **bash** shell, you can see how **exec** works by typing

```
exec ls -l
```

The **ls** command runs as usual, but when it is done, you are no longer logged in. The shell is replaced by **ls**, and when it finishes, it is as if your shell had finished.

How Does it All Get Started?

When the kernel is first loaded into memory, it initializes itself and any hardware that may be attached to the computer. Once the kernel is established enough to be able to run programs, it does. The first program is called “init”; its job is to function as the ancestor of all processes.

When **init** starts, it reads a file called **inittab**, usually located in **/etc**. This file tells **init** which programs should be run under which conditions. Not only does **init** run the start-up scripts that bring the rest of the system up, but **init** also takes care of shutting the system down.

One of the responsibilities of **init** is running the programs that let users log into the system. For a terminal (or virtual console), the two programs used are **getty** and **login**. **getty** is short for “get terminal”. A basic **getty** program opens the terminal device, initializes it, prints **login:** and waits for a user

Listing 1. autologin

```
#!/bin/bash
exec 0</dev/$1 1>/dev/$1 2>&1
cat /etc/issue
shift
exec $*
```

name to be entered. Modern **getty** programs (several are available for Linux) can do other things as well—if the terminal device is a (recent) modem, they can read status codes sent by the modem to tell if the call is voice or fax and handle the call appropriately. Most of the time, though, someone just wants to log in, so **getty** executes the **login** program, giving the user name to log in via the command line.

The **login** program then prompts the user for a password. If the password is wrong, **login** simply exits. The **init** process then notices that one of its children has exited and spawns another **getty** process on the terminal in question. If the password is good, **login** changes its process user ID to that of the user and executes the user's shell. At this point, the user can type commands. When the user exits by typing the shell's built-in **logout** command, the shell exits and **init** notices that its child has exited and spawns another **getty** on the terminal.

Why are two separate programs used to log in instead of just one? The answer is that doing it this way provides more flexibility. For example, **getty** doesn't have to execute **login**—it can execute a program to receive (or send) faxes, a PPP daemon to emulate a network connection over a serial line, or if you have a modem with “voicemail”, one of those phone tree programs that people hate so much (“press five to hear these options again”).

Similarly, **login** is sometimes needed without **getty**; for example, when a user logs in over a network, no terminal device is waiting. Instead, each new connection is handled by a program called **telnetd** that forks and executes a **login** process. **telnetd** remains to pass characters between the network and the new shell.

As a partial example of how the process works, Listing 1 shows an **autologin** replacement for **getty**. This replacement is meant for people who are tired of typing their user ID and password for the bazillionth time. You can boot Linux and have it drop straight into a couple of shells—sort of like DOS, but with virtual consoles.

To install **autologin**, copy it to the **/sbin** (system binaries) directory and type:

```
chmod +x /sbin/autologin
```

as root. Still as root, edit the **/etc/inittab** file and change the lines that look like this:

```
c1:12345:respawn:/sbin/getty 38400 tty1
```

to:

Linux Apprentice continued on page 95

while you are at it.

I guess most programs that decide to access /dev/tty to talk to a real user never consider the possibility that they do not have permission to do so.

In general, processes inherit an open file handle to /dev/tty from their parent as stdin, stdout and stderr. The process that opens them in the first place is **login** which is running as root, so it has no problems. Either that or the top-level program opens a particular real or virtual terminal using a device name which will have the permission bits set differently.

--

Adrian Pronk
apronk@sangacorp.com

Linux Journal and Red Hat

I have been a subscriber for about one year and I really like the articles. I am happy to see that Linux is stealing the spotlight from Microsoft.

I first got into Linux in 1996 after working with the HP-UX workstation on a job and realizing that you can connect to the Net with any OS, not just Macintosh, Windows 95 or Windows 3.x.

Red Hat is a good distribution and I have been running 5.0 for a year without any problems. I plan to get 5.2 soon.

For all you newbies, I recommend *Linux for Dummies*, *UNIX for Dummies* and *Teach Yourself Linux in 24 Hours*. These books include the operating system and some applications. I would recommend Red Hat to anyone.

--

Fred Nance
fnance@eclaim.com

VFS Error Messages

Regarding December's "Best of Technical Support", "VFS Error Messages": another reason for the failure to mount the root partition during boot is that the root file system was compiled as a module when the kernel was built, rather than compiled into the kernel itself. When the root file system is first mounted during the boot process, no modules are loaded, and modules cannot be loaded until the root file system is mounted. If the file system driver is a module, then the kernel cannot mount it—so it panics.

--

Rob Singleton
single@nortelnetworks.com

Real Life Business Story

As a system integration tool, Linux has allowed us to prepare custom network file servers which can do the following:

- Provide complete web-server services (Apache).
- Provide Internet connectivity for many users on the local LAN (IP-Masquerade).
- Provide file and printer services for Windows/DOS users (Samba).
- Provide file and printer services for Netware users (MARS_NWE, NCPFS).

- Provide complete internal/external e-mail services (Sendmail).
- Provide inexpensive terminals in both graphical and text-based platforms with the X Window System which can be connected in a variety of ways (Ethernet, serial, etc.).
- Provide complete point-to-point protocol (PPP) implementation for routing and other remote-oriented operations.
- Provide a fully scalable system that can grow with the company.

All of the above have been thoroughly tested and implemented. We could not be happier with the performance and continued development of this OS.

--

Larry Rivera
larrydog@coqui.net

Review of Learning the Bash Shell

In the review of *Learning the Bash Shell, Second Edition* (December 1998), Bob van der Poel points out that the book examples available by FTP are from the first edition and that the correct ones might be available by the time the review was printed. That should now be true, as I asked the publisher to correct this mistake in October.

--

Cameron Newham, Author
cam@sspl.demon.co.uk

Re: CIDR

In the December issue, there is a misprint in David A. Bandel's article "CIDR: A Prescription for Shortness of Address Space". On page 26, under the CIDR heading, second paragraph, it states:

For a Class C address, this default subnet is 24 bytes long, so putting all ones in the first 24 bytes and zeroes in the rest, we have 255.255.255.0.

I think this should instead read 24 bits, rather than bytes, since each octet is composed of 8 bits, which gives 4 total bytes. Just wanted to bring this to your attention. Great article!

--

Bob Cummings
bob@cter.eng.uab.edu


Mea culpa—you are 100% correct. Thank you for pointing out that little lapse of attention on my part. Guess I need to re-read everything thrice, because I read it twice and missed it both times.

--

David A. Bandel
dbandel@ix.netcom.com

Write us at ljeditor@ssc.com or send snail mail to Marjorie Richardson, Managing Editor, *Linux Journal*, P.O. Box 85867, Seattle, WA 98145-1876. All published letters are subject to editing. More Letters to the Editor can be found on our web site at <http://www.linuxjournal.com/issue##/>, where ## is the issue number of interest.

ACCOUNTING THAT FITS!



Accountflex fits a wide range of businesses by letting you select the features you want. Easily modifiable 4GL source lets you go for the perfect fit. Our 10 modules are portable, full featured, Informix compatible, and consistent in quality since 1990.



INFOFLEX, Inc.
840 Hinckley Road, Suite 107
Burlingame, CA 94010
650/697-6045 Fax 650/697-7696
www.infoflex.com
sales@infoflex.com

* Informix is copyrighted by Informix, Inc.

RS #037

is Your SQL a ROYALty PAIN?

ESQLFLEX provides ANSI standard SQL functionality from within "C" programs. **ESQLFLEX** is also a 100% compatible replacement for both Informix-ESQL/C and Informix-SE. Other **INFOFLEX** products include **QUERYFLEX**, a 'point and shoot' report writer, **INFOFLEX-4GL**, and **ACCOUNT-FLEX**. All products are *Informix compatible and available on UNIX, DOS, WINDOWS, and NT.



INFOFLEX, Inc.
840 Hinckley Road, Suite 107
Burlingame, CA 94010
650/697-6045 Fax 650/697-7696
www.infoflex.com
sales@infoflex.com

* Informix is copyrighted by Informix, Inc.

Prices start
at \$995
Developer
and \$125
per
Runtime.

RS #037

LinuxCAD 3D v2.0 WWW.LINUXCAD.COM

New version of popular Computer Aided Design and Drafting package. Includes 2D and 3D, line types, SHX text fonts, blocks and attributes. All major features of AutoCAD are implemented in LinuxCAD in a very compatible, user friendly way. Print and plot. Scripting language for parametrics.

Pricing:

for Linux Intel - only \$99
for LinuxPPC on PowerMac - only \$120
(847)8915971, (719)5229170,
(847)8917190

RS #282

BRIAN'S BOOKS

Specializing in computer books. Including:
Linux • UNIX • Networking • Internet • Java
Perl • Administration • Security • C++
and more! Discounted everyday!

Publishers such as

Addison Wesley • Hayden • MIS Press
O'Reilly • Prentice Hall • QUE • SAMS
SSC • Waite Group • Wiley • Wrox

888-755-2665 or 732-249-6492

books@davison.net www.briansbooks.com

RS #259

Linux Apprentice

continued from page 93

```
c1:12345:respawn:/sbin/autologin tty1 login -f myid
```

replacing `myid` with your own user ID. Red Hat installations typically do not have the letter `c` at the beginning of the line.

Be sure to leave some of the lines containing `getty` exactly as they are—if you do something wrong, you are going to need a way to log in to your system. On my own system, I change `c1` through `c3` and run three initial shells. Once the file is edited, reboot the system and all should work.

The first argument to `autologin` is the name of the terminal. The rest of the command line is used as the login command that does the work.

A Synopsis of autologin

The first line tells the kernel how to run this program, in this case by letting the bash shell interpret it. The first `exec` line is a Bourne shell trick that lets a shell script

change the source/destination of its standard input, standard output and standard error. We want to set file descriptors 0, 1 and 2 to refer to the terminal device as expected by `login` (and many other programs) when they run. The `cat` command displays the system's standard logon message. The `shift` command shifts the positional parameters to the shell script. Argument \$1 is deleted, argument \$2 becomes \$1, argument \$3 becomes \$2 and so on. The last line executes the rest of the command line as a program. In this case, the `login -f` option performs the normal login procedure, with the `-f` option telling `login` not to bother with passwords.¶



Andy Vaught is currently a Ph.D. candidate in computational physics at Arizona State University and has been running Linux since 1.1. He enjoys flying with the Civil Air Patrol as well as skiing. He can be reached at andy@maxwell.la.asu.edu.

HI-TECH USA

BEST SERVICES

Since 1990

LINUX/SOLARIS/SCO/NT SERVER SOLUTIONS


HIGH QUALITY

Custom Configuration

HI-TECH RS2002

\$4,318

- . Dual Pentium II Motherboard
- . W/Intel 440GX chipset AGP, W/RAID
- . 1X Pentium II 450 MHz CPU
- . 512MB DIMM SDRAM
- . Adaptec AAA-131 RAID Single Channel (Level 0 - Level 5)
- . Teac 1.44MB FDD
- . 10/100MBPS PCI Ethernet Card
- . 20GB(5x4.3 GB) ULW SCSI RM HD
- . 24X SCSI Int. CDROM
- . PCI W/2MB VGA CARD
- . ATX ServerTower W/ Dual 300W
- . Hot Swap Power Supply
- . Logitech PS/2 Mouse
- . Keytronic Enhanced Keyboard
- OPTIONS: Linux/Solaris/SCO/NT
- 400MHZ\$4,218 \$152 /MO
- 350MHZ\$3,928 \$138 /MO
- 333MHZ\$3,828 \$134 /MO
- 300MHZ\$3,798 \$133 /MO



**CALL for the
Intel Xeon 400/450
& Latest Price**

Since the Ads's price is always behind of the marketing price, please call our Sales Rep. or check our Web Site for latest price.

HI-TECH RMS3002


RACKMOUNT SERVER

\$2,328

- . SBC Board W/Pentium II/333MHz , Intel 440LX chipset
- . Passive Back Plane 9XPCI, 4XISA
- . 128MB DIMM SDRAM
- . UltraWide SCSI Controller
- . Teac 1.44MB FDD
- . 10/100MBPS PCI Ethernet Card
- . 9 GB ULTRA Wide SCSI HD
- . 32X SCSI Int. CDROM
- . ATI PCI W/2MB
- . ATX Rackmount Case W/ 300W P/S
- . Logitech PS/2 Mouse
- . Keytronic Enhanced Keyboard

OPTIONS: Linux/Solaris/SCO/NT

- 300MHZ\$2,288 \$ 81/MO
- 266MHZ\$2,268 \$ 80/MO



HITECH LS2001

Linux Workstation

\$888

BUSINESS LEASE \$31/MO

- . Intel Chipset M/B
- . Pentium 233MHz MMX CPU
- . 64MB EDO RAM
- . 32XEIDE CDROM
- . Diamond 4MB Video Card
- . 4.3GB Ultra EIDE HD
- . 56.6KBPS INT. Fax/Modem
- . Keyboard, Mouse, 1.44MB FDD
- . MidTower Case w/230W PS
- . Win 98 w/CD, Manual
- . 15" SVGA NI .28mm Monitor
- . RedHat Linux (pre-loaded)
- . Call for the other Models and OSs (Solaris/SCO/NT)




HI-TECH WIN500AX

Sun AXI PCI SERVER

\$3,798

BUSINESS LEASE \$139/MO

- . Sun SP ARCengine Ultra Axi w/300/MO
- . 64MB EDO DIMM .Ultra Wide SCSI Controller
- . 1.44MB FDD .4.3GB .Ultra Wide SCSI Int. HD
- . 24X SCSI INT. CDROM .ATI PCI VGA w/4MB
- . Built-in 10/100 MHz Ethernet .Mouse .Keyboard
- . Solaris V.2.6 Server



ULTRASPARC[®]
DRIVEN

Software:

- . Solaris V.2.6 Desktop(Int/SP)\$295
- . Solaris V.2.6 Server W/5 user.....\$495
- . Solaris V.2.6 Server(unlimited).....\$1525
- . SCO OpenServer DT\$615
- . SCO Host Server w/5 user\$554
- . SCO Enterprise w/5 user\$1065
- . Slackware Linux\$25
- . RedHat Linux V.5.1\$45
- . PC to UNIX for Win 95/NT\$195
- . Novell Netware 4.11 5/10 \$695/\$1295
- . PC Anywhere Remote&Host 8.0.....\$135
- . WinFax Pro 4.1 Net. W/25 user..\$1355
- And more....


Hardware:

Cyclade:

- . 4 Ports, DB25 Ext. Card \$228
- . 8 Ports, RJ45/DB25, Int./Ext... \$238/\$448
- . 16 Ports, DB25, ISA Ext. \$588
- . PR3000 W/2S at T1/1:1 SP \$1878
- . PR3000 Term. Server 16 Port..... \$2458
- . Router+DSU/CSU Link-56K..... \$998

Ascend:

- . Pipeline 50 Base w/NT1\$618
- . Pipeline 50 Router, S/W 56/CSU/DSU\$1238
- . Pipeline 75 ISDN Bri/Rout w/Comp\$698
- . Pipeline 85 ISDN Bri/rout w/Sec. F/W\$1238



Warranty:

One Year Parts & 5 Year Labor



HI-TECH USA

1562 Centre Pointe DR.
Milpitas, CA 95035

Tel: (888)868-8778

(408)946-6886

Fax: (408)262-8868

Email: sales@hitech-usa.com

Web: http://www.hitech-usa.com

CORPORATE/GOVERNMENTAL/EDUCATIONAL INSTITUTIONS WELCOME

LEASING & FINANCING AVAILABLE

We do sell relative components and software.

Business Hours: Mon. - Fri. 8AM to 6PM(PST)

Solaris is a trademark of Sun Microsystems, Inc. SPARC is a trademark of SPARC international. All other trademarks are property of their respective owners.

IF YOU'RE READY, WE'RE READY


11950 N. Highway 89
Flagstaff, AZ 86004

InfoMagic

1-800-800-6613
Tel +1-520-526-9565
Fax +1-520-526-9573
orders@infomagic.com

LINUX Developer's Resource 6 CD Set

This 6-CD set contains the major distributions of Linux. Current releases of the following distributions are generally included:



- RedHat for Intel
- SuSE
- Slackware®
- Debian

Also includes a "QuickStart" guide that describes the installation and configuration process for the InfoMagic release, **not** the "generic" installation guide. In addition, it comes with complete online documentation from the Linux Documentation Project!

InfoMagic makes contributions to the Free Software Foundation, XFree86 Project, Linux Documentation Project and the Linux Grant Program.

LINUX Developer's Resource CD Subscription Service

Purchase the first release of our Linux Developers Resource CD set at the full price and request the subscription service to automatically receive all future updates at a discounted price plus shipping. Subscribers are always the first ones to receive new CD releases hot off the press! The newest release will be billed to your credit card at the time each update ships. Updates will ship automatically until you cancel.

LINUX Archives




Linux Archives is the complete mirror from the most popular archive sites! Includes thousands of files for adding on and extending the power of whatever linux distribution you are running. The following archives contain a tremendous wealth of peripheral software for Linux such as internet utilities, multimedia tools, networking software, development tools, compilers, drivers,

communication tools and more! The following sites are included:

- SUNSITE Archive (sunsite.unc.edu)
- TSX-11 Archive (tsx-11.mit.edu)
- GNU Archive (prep.ai.mit.edu)
- KERNEL Archive (ftp.kernel.org)
- KDE Archive (ftp.kde.org)
- XFree86 (ftp.xfree86.org) (Intel binaries and full source)
- XFree86 contrib (ftp.xfree86.org)
- LessTif (ftp.lesstif.org)
- Apache (ftp.apache.org)
- GNOME (ftp.gnome.org)

The Linux Archive subscription service is also available! Purchase the first version at full price and we will automatically ship you the new releases. A subscription requires a credit card and will continue until you cancel.

LINUX Installation & Getting Start Guide by Matt Welsh




This 245 pg manual is the complete guide to installing and using the Linux operating system. The purpose of this book is to get users up & running with ease by providing the basics needed in order to run Linux & UNIX implementations. Chapters include an introduction to Linux, history, software features, design philosophy, hardware requirements & sources for other documentation.

LINUX T-SHIRTS!

Top quality 100% Cotton T-shirts. Front has the InfoMagic "ELFOS" character with LINUX in big bold letters. On the back: "LINUX: The Operating System with an Attitude". Pre-shrunk sizes: M, L, XL, XXL.



LINUX TOOLBOX



The perfect package for beginners! InfoMagic has bundled our best Linux CD with the most in depth manual available for Linux in one complete package!


The LINUX DEVELOPERS RESOURCE 6 CD-ROM Set includes complete materials to install and run LINUX. Current releases of the following distributions and helpful programs are included:

- RedHat
- SuSE
- XFree86
- Slackware
- Debian
- Apache

The CD Set includes a simple "QuickStart" install booklet with basic setup instructions. Complete on-line documentation from the Linux Documentation Project along with the "Installation & Getting Started Guide" and the "Network Administrators Guide".

"RUNNING LINUX" by Matt Welsh & Lar Kaufmann — This book published by O'Reilly & Associates covers everything you need to know in order to understand, install and run Linux. Includes chapters on UNIX basics for Linux users, system administration and maintenance, tools available for the system, tools for interfacing with MS-DOS, and Network configuration and administration. This is the most comprehensive LINUX package available!

MOO-TIFF (for Linux) 2 CD Set




Finally! An affordable version of the popular 100% OSF/MOTIF-compatible Graphical User Interface for PC's running LINUX! Includes complete runtime and development environment, provides everything needed to run mwm and develop 100% OSF/Motif-compatible applications. Derived & built from the original Motif 2.0 sources from the Open Software Foundation. Requires

Linux 1.0 (or higher) & XFree86 2.x (or higher). **Packages include:** Static Libraries (libXm, LibMrm & libUil), Header & Include Files, Shared Library (libXm, LibMrm). On-line user's guide! **Chapters include:** Introduction to OSF/Motif, Using the Motif Window Manager, Using Customization Features, Interacting with Motif Aps, Customizing the Window Manager, Customizing Window Apps. Demos include: Clipboard-cut & paste between aps, hel-lomotif - button-press demo, periodic-a table of Motif Widgets, textedit-a text editor, xmpiano-a music editor example, xmsamplers- Motif demos, DragAndDrop - sample Drag & Drop aps (requires color screen).

INFOMAGIC WORKGROUP SERVER

Using proven software, this package provides high-performance networking, file, printing and fax services to PC and Mac clients using the Linux operating system. Already in use worldwide, the InfoMagic Workgroup Server provides simple graphical tools for system administration and set-up. Features include:



PC Services: works with Windows for Workgroups, Windows 95 & NT, OS/2 or DOS, no server-specific software on the client machine and no license charges, mounts server file directories using the usual mechanisms, connects to server print queues and print as if printing to a local printer, allows easy sending of faxes to the fax printer

Macintosh Services: works with AppleTalk client software, mounts server disk volumes using the Chooser, prints as if printing to any other printer, faxes simply by printing to the appropriate printer

Firewalling capability: allows PCs onto the internet, without fees for extra ISP connections, allows everyone access via a single PPP modem or ISDN connection, allows PCs on the internet, while restricting employee browsing, outside sites for email, ftp data, WWW data, and more, access on the Server can be enabled or disabled

WWW services: saves PC or Mac files into a shared drive or network folder, combine the WWW server, Macintosh and PC services, high performance apache WWW server.

Call, Fax or send e-mail
MC, VISA and AMEX accepted

Visit our website at <http://www.infomagic.com>

METRO 3D EXTREME

HARDWARE - ACCELERATED OpenGL®

METRO 3D EXTREME

Hardware accelerated OpenGL®

3D Hardware Support
High-Performance Accelerated
Implementation of OpenGL
OpenGL Conformance
Test Certified
Includes a Free Copy of METRO-X
Available for:
Linux/x86, Linux/Alpha, & FreeBSD

METROLINK MOTIF COMPLETE!

The ULTIMATE Motif for Linux

3 Versions of Motif on 1 CD
Multiple Development Environments
Mix & Match Modules with
Graphical Installation
Includes Motif Ver. 1.2, 2.0, & 2.1
(Both glibc and libc5)
Available for:
Linux/x86, Linux/Alpha, & FreeBSD

METRO-X

A reliable, easy-to-install, high
performance X server for Linux

Graphical Configuration Utility
Touch Screen Support
Multi - Screen Support
3D Input Device Support
Robust and High - Performance
X Server
Available for:
Linux/x86, Linux/Alpha, FreeBSD,
BSDI, Lynx, Solaris/x86, and QNX



OpenGL is a registered trademark of Silicon Graphics Incorporated

4711 Powerline Road, Ft Lauderdale, FL 33309 (954) 938-0283 FAX (954) 938-1982
email: sales@metrolink.com http://www.metrolink.com